

IBM Storage Protect
8.1.24

Multisite Disk Solution Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 143.](#)

Edition notice

This edition applies to version 8, release 1, modification 24 of IBM Storage Protect (product numbers 5725-W98, 5725-W99, 5725-X15), and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 1993, 2024.**

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

Contents

About this publication.....	vii
Who should read this guide.....	vii
Publications	vii
What's new.....	ix
Part 1. Planning.....	1
Selecting a system size.....	4
Site Planning for multi-target replication solution.....	4
System requirements for a multisite disk solution.....	6
Hardware requirements.....	6
Software requirements.....	8
Planning worksheets.....	10
Planning for storage.....	21
Planning the storage arrays.....	22
Planning for security.....	23
Planning for administrator roles.....	23
Planning for secure communications.....	24
Planning for storage of encrypted data.....	25
Planning firewall access.....	25
Part 2. Implementation	27
Setting up the system.....	28
Configuring the storage hardware.....	28
Installing the server operating system.....	28
Installing on AIX systems.....	28
Installing on Linux systems.....	30
Installing on Windows systems.....	35
Configuring multipath I/O.....	35
AIX systems.....	35
Linux systems.....	36
Windows systems.....	37
Creating the user ID and directories for the server instance.....	38
Preparing file systems for the server.....	41
AIX systems.....	41
Linux systems.....	42
Windows systems.....	43
Installing the server and Operations Center.....	44
Installing on AIX and Linux systems.....	44
Installing prerequisite RPM files for the graphical wizard.....	45
Installing on Windows systems.....	45
Configuring the server and the Operations Center.....	46
Configuring the server instance.....	46
Installing the backup-archive client.....	47
Setting options for the server.....	48
Configuring secure communications with Transport Layer Security.....	49
Configuring the Operations Center.....	49
Securing communications between the Operations Center and the hub server.....	50
Registering the product license.....	52
Configuring data deduplication.....	52

Defining data retention rules for your business.....	53
Defining schedules for server maintenance activities.....	53
Defining client schedules.....	56
Installing and configuring backup-archive clients.....	56
Registering and assigning clients to schedules.....	56
Installing the client management service.....	57
Verifying that the client management service is installed correctly.....	58
Configuring the Operations Center to use the client management service.....	59
Configuring the second server.....	60
Configuring SSL communications between the hub server and a spoke server.....	60
Adding the second server as a spoke.....	62
Enabling replication.....	62
Completing the implementation.....	62

Part 3. Monitoring..... 65

Daily checklist.....	65
Periodic checklist.....	77
Verifying license compliance.....	82
Tracking system status by using email reports.....	84

Part 4. Managing..... 87

Managing the Operations Center.....	87
Adding and removing spoke servers.....	87
Adding a spoke server.....	87
Removing a spoke server.....	88
Starting and stopping the web server.....	88
Restarting the initial configuration wizard.....	89
Changing the hub server.....	90
Restoring the configuration to the preconfiguration state.....	90
Protecting applications, virtual machines, and systems.....	92
Adding clients.....	92
Selecting the client software and planning the installation.....	93
Specifying rules for backing up and archiving client data.....	94
Scheduling backup and archive operations.....	97
Registering clients.....	98
Installing and configuring clients.....	99
Managing client operations.....	103
Evaluating errors in client error logs.....	103
Stopping and restarting the client acceptor.....	104
Resetting passwords.....	105
Modifying the scope of a client backup.....	106
Managing client upgrades.....	107
Decommissioning a client node.....	108
Deactivating data to free storage space.....	110
Managing data storage.....	111
Auditing a storage pool container.....	111
Managing inventory capacity.....	111
Managing memory and processor usage.....	113
Tuning scheduled activities.....	114
Moving clients.....	115
Managing replication.....	116
Replication compatibility.....	116
Replicating client data by using replication storage rules.....	117
Enabling node replication.....	118
Protecting data in directory-container storage pools.....	118
Modifying replication settings.....	120
Setting different retention policies.....	121

Replicating client node data after a database restore.....	122
Securing the server.....	123
Security concepts.....	123
Managing administrators.....	126
Changing password requirements.....	126
Securing IBM Storage Protect on the system.....	128
Restricting user access to the server.....	128
Limiting access through port restrictions.....	129
Stopping and starting the server.....	130
Stopping the server.....	130
Starting the server for maintenance or reconfiguration tasks.....	131
Planning to upgrade the server.....	132
Preparing for an outage.....	132
Implementing a disaster recovery plan.....	133
Recovery drills.....	133
Recovering from data loss or system outages.....	134
Restoring the database.....	136
Repairing storage pools.....	137
Synchronization of source and target replication servers after role reversal.....	138
Appendix A. Accessibility.....	141
Notices.....	143
Glossary.....	147
Index.....	149

About this publication

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

Who should read this guide

This guide is intended for anyone who is registered as an administrator for IBM Storage Protect. A single administrator can manage IBM Storage Protect, 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 or 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 IBM Storage Protect product family includes IBM Storage Protect Plus, IBM Storage Protect for Virtual Environments, IBM Storage Protect for Databases, and several other storage management products from IBM®.

To view IBM product documentation, see [IBM Documentation](#).

What's new in this release

This release of IBM Storage Protect introduces new features and updates.

For a list of new features and updates in this release, see the following topics:

- [What's new for Server components](#)
- [What's new for Client components](#)

If changes were made in the documentation, they are indicated by a vertical bar (|) in the margin.

Part 1. Implementation methods and roadmap plan for a multisite disk data protection solution

Plan for a multi-target replication solution with servers configured at two sites that use data deduplication and replication.

Implementation methods

You can configure servers for a multi-target replication solution in the following ways:

Configure servers by using the Operations Center and administrative commands

You can 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 IBM Storage Protect commands. For more information, see [“Planning a roadmap” on page 1](#).

Configure servers by using automated scripts

For detailed guidance on configuration with specific IBM Storwize® storage systems, and by using automated scripts to configure each server, see [IBM Storage Protect Blueprints](#).

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 Storage Scale technology, is included.

Planning a roadmap

Plan for a multisite disk solution by reviewing the architecture layout in the following figure and then completing the roadmap tasks that follow the diagram.

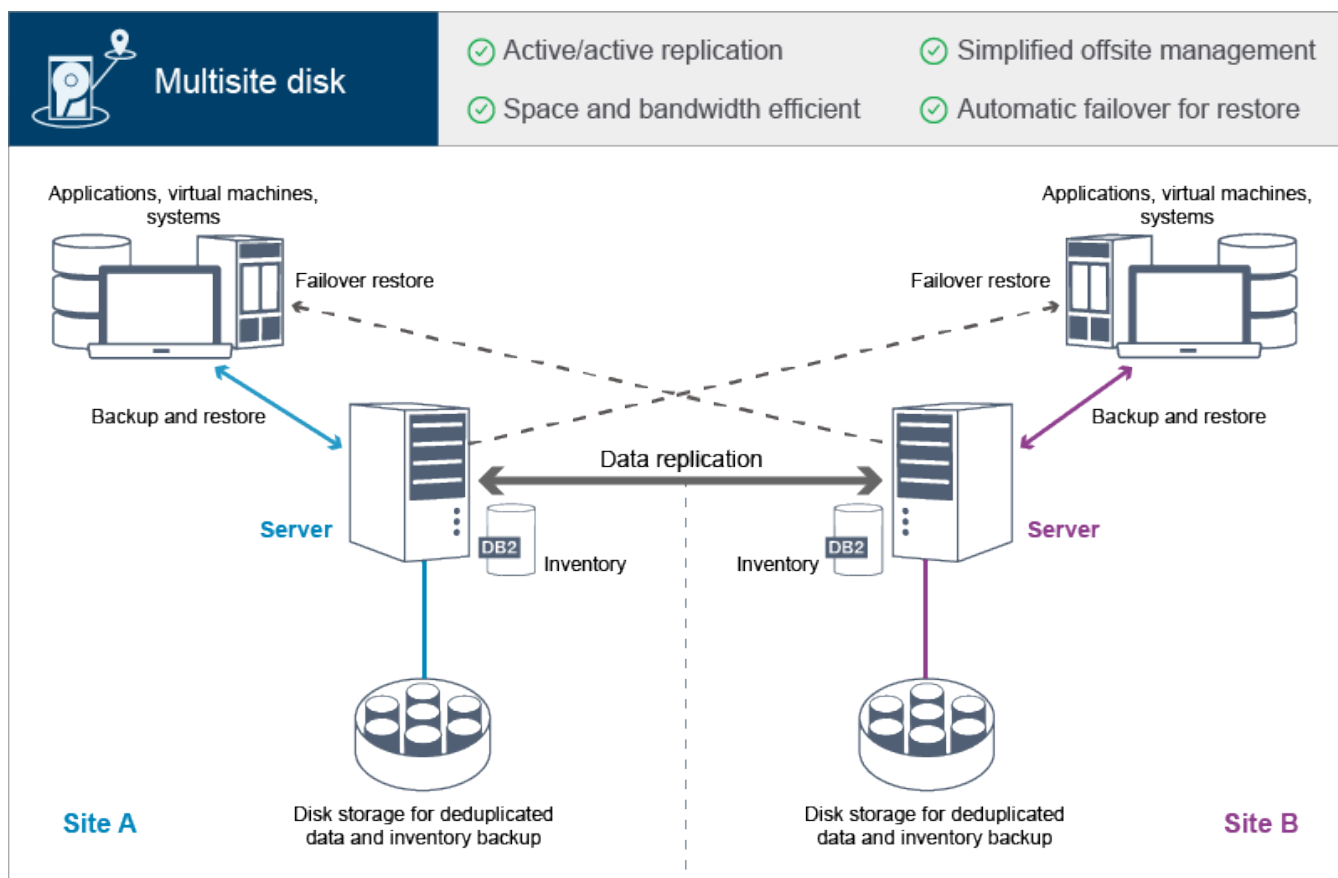


Figure 1. Multisite disk solution

Plan for a multi-target replication solution by reviewing the architecture layout in the following figure and then completing the roadmap tasks that follow the diagram.

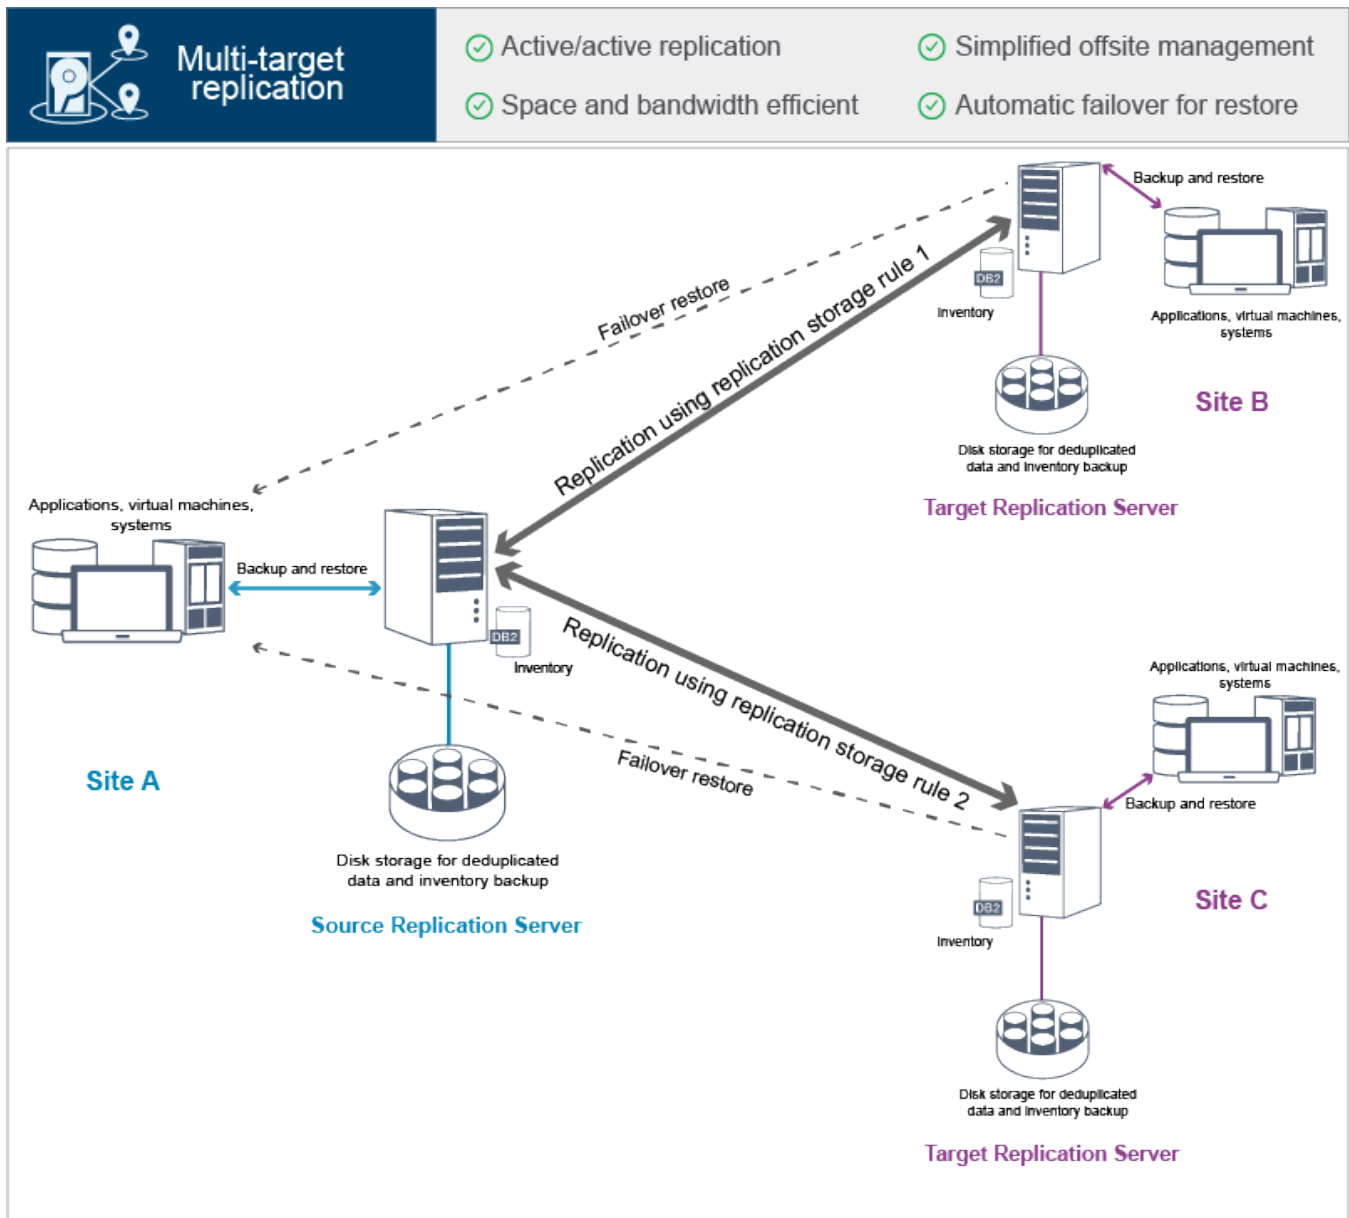


Figure 2. Multi-target replication solution

Tip: IBM Storage Protect 8.1.13 introduces a feature for replicating data by defining replication storage rules and associated subrules. This feature streamlines the configuration process and supports fine-tuning of replication operations. In addition, the feature supports enhanced protection of data in directory-container storage pools. If you implemented replication by using the **REPLICATE NODE** command, consider a transition to replication storage rules and subrules.

Roadmap tasks

The following steps are required to plan for a multisite disk or a multi-target replication 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.

Selecting a system size

Select the size of the IBM Storage Protect 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 the 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.

<i>Table 1. Determining the size of the server</i>			
Total managed data	Daily amount of new data to back up with one replication copy	Daily amount of new data to back up with two replication copies	Required server size
10 TB - 40 TB	Up to 1 TB per day	Up to 0.6 TB per day	Extra Small
60 TB - 240 TB	Up to 10 TB per day	Up to 6 TB per day	Small
360 TB - 1440 TB	10 - 30 TB per day	6 - 18 TB per day	Medium
1000 TB - 4000 TB	30 - 100 TB per day	18 - 60 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 IBM Storage Protect for Virtual Environments. Workloads that consist of objects that are smaller than 128 KB might not be able to achieve these daily limits.

Remember: If you are planning to create two replication copies of the backup data, you will need to consider it while selecting the size of the server. The daily amount of backup data has to be decreased to reduce the amount of time required to back up data. This is done to compensate for the additional time needed to create the second replication copy.

Site Planning for multi-target replication solution

Review use cases and evaluate the factors to provide the most effective data protection for the multi-target replication solution for IBM Storage Protect.

Use cases

With the multi-target replication solution, you can create a copy of backup data from a source replication server to two target replication servers.

Tip: Avoid conflicts in managing administrative IDs and client option sets by identifying the IDs and option sets that are replicated to the target replication server and the IDs and option sets that are managed in an enterprise configuration. You cannot define two administrative user IDs for the same registered node.

Although your company might benefit from a multi-target replication solution for various reasons, the most common reasons to use a multi-target replication 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 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 server at Site B to recover systems. The directory container storage pool at site A can be repaired from the copies at either site B or site C. In addition, client systems at site A can restore data from the server at either site B or site C.

Replication from the source replication server (primary site) to the two target replication server (disaster recovery sites)

In this scenario, data that is backed up to the source replication server, Site A, is replicated to two target replication servers for disaster recovery, Site B and Site C. If a disaster occurs at Site A, such as failure of the server, you can use either server at Site B or server at Site C to recover systems. Alternatively, you can use the server at Site A to restore primary storage pool data at Site B or Site C, such as after a disk storage failure at either of the target replication servers.

Mutual replication at two active sites

In this scenario, local data at each site is backed up by the 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 servers to the primary site

In this scenario, you configure remote servers that are relatively small to replicate data that is backed up to a larger 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 multi-target replication 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 \(V7.1.1\)](#).

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 activity to only a portion of the bandwidth. There are no controls in IBM Storage Protect 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 replication 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 replication servers on the same site.

Daily data ingestion

For the multi-target replication solution, the daily data ingestion and total data retention must be within the capacity of the configurations. For example, a large configuration has a data ingestion capacity of up to 60 TB per day, including node replication. In cases where the backup requirements exceed the capacity of a single server, you can configure a solution that uses multiple servers to achieve the required capacity.

Server configuration

The server configuration must meet or exceed the requirements for the multi-target replication solution.

Single replica of backed-up data

A data replication 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 replication server. If you require multiple backed-up data copies, you can consider the multi-target replication solution to create copies of data to two target replication servers.

Related reference

[System requirements for a multisite disk solution](#)

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

System requirements for a multisite disk solution

After you select the IBM Storage Protect 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

[IBM Storage Protect Supported Operating Systems](#)

Hardware requirements

Hardware requirements for your IBM Storage Protect solution 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 [Selecting a system size](#).

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.

Use the information in the following table as a starting point. For the most up-to-date information about hardware requirements and specifications for the server and storage, see [IBM Storage Protect Blueprints](#).

Hardware component	Small system	Medium system	Large system
Server processor	<div><div>AIX</div>6 processor cores, 3.42 GHz or faster</div> <div><div>Linux</div><div>Windows</div>16 processor cores, 1.7 GHz or faster</div>	<div><div>AIX</div>10 processor cores, 3.42 GHz or faster</div> <div><div>Linux</div><div>Windows</div>20 processor cores, 2.2 GHz or faster</div>	<div><div>AIX</div>20 processor cores, 3.42 GHz</div> <div><div>Linux</div><div>Windows</div>44 processor cores, 2.2 GHz or faster</div>

Hardware component	Small system	Medium system	Large system
Server memory	64 GB RAM	128 GB RAM	256 GB RAM
Network	<ul style="list-style-type: none"> • 10 GB Ethernet (1 port) • 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> • 10 GB Ethernet (2 ports) • 8 GB Fibre Channel adapter (2 ports) 	<ul style="list-style-type: none"> • 10 GB Ethernet (4 ports) • 8 GB Fibre Channel adapter (4 ports)
Storage	<ul style="list-style-type: none"> • 1.45 TB SSD disks for the database, plus space for Operations Center records • 67 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> • 2.53 TB SSD disks for the database, plus space for Operations Center records • 207.9 TB deduplicated directory-container storage pool 	<ul style="list-style-type: none"> • 6.54 TB SSD disks for the database, plus space for Operations Center records • 1049.67 TB deduplicated directory-container storage pool

Implementing the correct processor core technology

You must use the correct type of processor core technology for the server processor. For information about the type of processor core technology, see [IBM Storage Protect Blueprints](#).

Estimating database space requirements for the Operations Center

Hardware requirements for the Operations Center are included in the preceding table, except for the database and archive log space (inventory) that the Operations Center uses to hold records for managed clients.

If you do not plan to install the Operations Center on the same system as the server, you can estimate system requirements separately. To calculate system requirements for the Operations Center, see the system requirements calculator in [technote 1641684](#).

Managing the Operations Center on the server is a workload that requires extra space for database operations. The amount of space depends on the number of clients that are monitored on a server. Review the following guidelines to 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 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:

$$7 \times 1.2 \text{ GB} = 8.4 \text{ GB}$$

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 each spoke server in the example, the archive log space that is used over 24 hours is approximately 16 GB. These estimates are based on the default status collection interval of 5 minutes. If you reduce 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 are 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 IBM Storage Protect 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	<p>IBM AIX 7.2</p> <ul style="list-style-type: none">• AIX 7.2 TL4 with SP3, TL5 with SP2, or later AIX 7.2 levels <p>IBM AIX 7.3</p> <ul style="list-style-type: none">• AIX 7.3. TL0 with SP1 or later AIX 7.3 levels <p>For more information about operating system requirements, see the IBM Storage Protect installation information.</p>
Gunzip utility	<p>The gunzip utility must be available on your system before you install or upgrade the IBM Storage Protect server. Ensure that the gunzip utility is installed and the path to it is set in the PATH environment variable.</p>
File system type	<p>JFS2 file systems</p> <p>AIX systems can cache a large amount of file system data, which can reduce memory that is required for server and IBM Db2 processes. To avoid paging with the AIX server, use the <code>rbw</code> mount option for the JFS2 file system. Less memory is used for the file system cache and more is available for IBM Storage Protect.</p> <p>Do not use the file system mount options, Concurrent I/O (CIO), and Direct I/O (DIO), for file systems that contain the IBM Storage Protect database, logs, or storage pool volumes. These options can cause performance degradation of many server operations. IBM Storage Protect and Db2 can still use DIO where it is beneficial to do so, but IBM Storage Protect does not require the mount options to selectively take advantage of these techniques.</p>
Other software	<p>Korn Shell (ksh)</p>

Linux systems

Type of software	Minimum software requirements
Operating system	<p>Red Hat® Enterprise Linux® 9.1 or later RHEL 9 levels</p> <p>Note: IBM Storage Protect server supports RHEL 9 levels on Linux x86_64 platform only.</p> <p>Red Hat Enterprise Linux 8.1 or later RHEL 8 levels</p> <p>Red Hat Enterprise Linux 7.6 or later RHEL 7 levels</p> <p>SUSE Linux Enterprise Server 15, SP1 or later SLES 15 levels</p> <p>SUSE Linux Enterprise Server 12, SP4 or later SLES 12 levels</p> <p>Ubuntu Server LTS 18.04 or later Ubuntu Server LTS 18.04 levels</p> <p>Ubuntu Server LTS 20.04 or later Ubuntu Server LTS 20.04 levels</p> <p>Note: The list of Linux versions supported only applies to xLinux and Linux on Power Systems. For Linux on System z it is limited to the following levels:</p> <p>Red Hat Enterprise Linux 7.6 or later RHEL 7 levels</p> <p>Red Hat Enterprise Linux 8.4 or later RHEL 8 levels</p> <p>SUSE Linux Enterprise Server 12, SP4 or later SLES 12 levels</p>
Libraries	<p>GNU C libraries, Version 2.3.3-98.38 or later that is installed on the IBM Storage Protect system.</p> <p>For Red Hat Enterprise Linux Servers:</p> <ul style="list-style-type: none"> • libaio • libstdc++.so.6 (32-bit and 64-bit packages are required) • numactl.x86_64 <p>For Red Hat Enterprise Linux (RHEL 8), also add this library:</p> <ul style="list-style-type: none"> • libnuma.so.1 <p>For SUSE Linux Enterprise Servers:</p> <ul style="list-style-type: none"> • libaio • libstdc++.so.6 at version 4.3 or later (32-bit and 64-bit packages are required) • libnuma.so.1 <p>For Ubuntu LTS Servers:</p> <ul style="list-style-type: none"> • libaio1 • libnuma.so.1
File system type	<p>Format database-related file systems with ext3 or ext4.</p> <p>For storage pool-related file systems, use XFS.</p>
Other software	Korn Shell (ksh)

Windows systems

Type of software	Minimum software requirements
Operating system	Microsoft Windows Server 2019 Microsoft Windows Server 2016
File system type	NTFS
Other software	Windows 2019 or Windows 2016 with .NET Framework 3.5 is installed and enabled. The following User Account Control policies must be disabled: <ul style="list-style-type: none">• User Account Control: Admin Approval Mode for the Built-in Administrator account• User Account Control: Run all administrators in Admin Approval Mode

Related information

[Setting AIX network options](#)

Planning worksheets

Use the planning worksheets to record values that you use to set up your system and configure the IBM Storage Protect server. Use the default values that are listed in the worksheets.

Each worksheet helps you prepare for different parts of the system configuration by using the default 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 IBM Storage Protect during system setup. All directories that you create for the server must be empty.

Server configuration

Use the configuration worksheets when you configure the server. Default values are suggested 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 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.

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

Item	Default value	Your value	Minimum directory size	Notes
Directory for the server instance	/home/tsminst1/tsminst1		100 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	/		5 GB	
Directory for server installation	/usr		5 GB	
Directory for server installation	/var		5 GB	
Directory for server installation	/tmp		5 GB	
Directory for server installation	/opt		10 GB	
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Windows Extra small: 30 GB • Small and medium: 140 GB • Large: 550 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		<ul style="list-style-type: none"> • Windows Extra small: 250 GB • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	

Table 2. Worksheet for preconfiguration of an AIX 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: <ul style="list-style-type: none"> • Windows Extra small: At least 200 GB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 4 file systems • Medium: At least 4 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: <ul style="list-style-type: none"> • Windows Extra small: At least 10 TB • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Windows Extra small: At least 2 file systems • Small: At least 2 file systems • Medium: At least 10 file systems • Large: At least 30 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: <ul style="list-style-type: none"> • Windows Extra small: At least 1 TB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 2 file systems • Medium: At least 3 file systems • Large: At least 3 file systems The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Table 3. Worksheet for IBM Storage Protect 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 10 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 hostname.		

<i>Table 3. Worksheet for IBM Storage Protect configuration (continued)</i>			
Item	Default value	Your value	Notes
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
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		<p>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.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Linux

<i>Table 4. Worksheet for preconfiguration of a Linux server system</i>				
Item	Default value	Your value	Minimum directory size	Notes
TCP/IP port address for communications with the server	1500		Not applicable	<p>Ensure that this port is available when you install and configure the operating system.</p> <p>The port number can be a number in the range 1024 - 32767.</p>

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

Item	Default value	Your value	Minimum directory size	Notes
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 17.
Directory for the active log	/tsminst1/TSMalog		<ul style="list-style-type: none"> • Windows Extra small: 30 GB • Small and medium: 140 GB • Large: 550 GB 	
Directory for the archive log	/tsminst1/ TSMarchlog		<ul style="list-style-type: none"> • Windows Extra small: 250 GB • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	
Directories for the database	/tsminst1/ TSMdbspace00 /tsminst1/ TSMdbspace01 /tsminst1/ TSMdbspace02 /tsminst1/ TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 200 GB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems

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

Item	Default value	Your value	Minimum directory size	Notes
Directories for storage	/tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 10 TB • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Windows Extra small: At least 2 file systems • Small: At least 2 file systems • Medium: At least 10 file systems • Large: At least 30 file systems
Directories for database backup	/tsminst1/TSMbkup00 /tsminst1/TSMbkup01 /tsminst1/TSMbkup02 /tsminst1/TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 1 TB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 2 file systems • Medium: At least 3 file systems • Large: At least 3 file systems The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Table 5. Worksheet for IBM Storage Protect 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.
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 hostname.		
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
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		<p>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.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Windows

Because many volumes are created for the 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. The volume 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 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 20 .
Directory for the active log	C:\tsminst1\TSMalog		<ul style="list-style-type: none"> • Windows Extra small: 30 GB • Small and medium: 140 GB • Large: 550 GB 	
Directory for the archive log	C:\tsminst1\TSMarch log		<ul style="list-style-type: none"> • Windows Extra small: 250 GB • Small: 1 TB • Medium: 2 TB • Large: 4 TB 	

Table 6. Worksheet for preconfiguration of a Windows server system (continued)				
Item	Default value	Your value	Minimum directory size	Notes
Directories for the database	C:\tsminst1\TSMdbspace00 C:\tsminst1\TSMdbspace01 C:\tsminst1\TSMdbspace02 C:\tsminst1\TSMdbspace03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 200 GB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 4 file systems • Medium: At least 4 file systems • Large: At least 8 file systems
Directories for storage	C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 10 TB • Small: At least 38 TB • Medium: At least 180 TB • Large: At least 500 TB 	Create a minimum number of file systems for storage, depending on the size of your system: <ul style="list-style-type: none"> • Windows Extra small: At least 2 file systems • Small: At least 2 file systems • Medium: At least 10 file systems • Large: At least 30 file systems

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

Item	Default value	Your value	Minimum directory size	Notes
Directories for database backup	C:\tsminst1\TSMbkup00 C:\tsminst1\TSMbkup01 C:\tsminst1\TSMbkup02 C:\tsminst1\TSMbkup03		Minimum total space for all directories: <ul style="list-style-type: none"> • Windows Extra small: At least 1 TB • 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 system: <ul style="list-style-type: none"> • Windows Extra small: At least 1 file system • Small: At least 2 file systems • Medium: At least 3 file systems • Large: At least 3 file systems The first database backup directory is also used for the archive log failover directory and a second copy of the volume history and device configuration files.

Table 7. Worksheet for IBM Storage Protect 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 18 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 hostname.		

Table 7. Worksheet for IBM Storage Protect configuration (continued)			
Item	Default value	Your value	Notes
Server password	passw0rd		Select a different value for the server password than the default. Ensure that you record this value in a secure location.
Administrator ID: user ID for the server instance	admin		
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		<p>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.</p> <p>Schedules for server maintenance operations are defined to begin 10 hours after the start of the client backup window.</p>

Planning for storage

Choose the most effective storage technology for IBM Storage Protect 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 IBM Storage Protect. For general guidance on selecting the appropriate storage hardware and set up for your solution, review the following guidelines.

Database and active log

- Use a fast disk for the IBM Storage Protect 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

Related information

[Storage system requirements and reducing the risk of data corruption](#)

Planning the storage arrays

Prepare for disk storage configuration by planning for RAID arrays and volumes, according to the size of your IBM Storage Protect system.

You design storage arrays with size and performance characteristics that are suitable for one of the IBM Storage Protect 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. If you use a smaller disk storage pool to stage data, you might need to use faster disks to manage the daily workload performance for ingesting and migrating data.

Table 8 on page 22 and Table 9 on page 23 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 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: <div>Quantity x Usable size</div>
Suggested storage group and array names	Preferred name to use for MDisks and MDisk groups.

Table 8. Components of storage group configuration (continued)	
Component	Details
Usage	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 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	Server component that uses part of the physical disk.

Examples

Configuration examples for storage groups and volumes are available at the following link: [Examples of 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.

Planning for security

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

For guidelines about protecting your storage environment against ransomware attacks, and recovering your storage environment if an attack occurs, see [Protecting the storage environment against ransomware](#).

Planning for administrator roles

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

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

System

Administrators with 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 on page 24](#) provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

<i>Table 10. Scenarios for administrator roles</i>	
Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	<ul style="list-style-type: none">• System authority: 1 administrator ID
An administrator for multiple servers also manages the overall system. Several other administrators manage their own storage pools.	<ul style="list-style-type: none">• 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 IBM Storage Protect servers.	<ul style="list-style-type: none">• 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 IBM Storage Protect 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. To improve system performance, use TLS for authentication without encrypting object data. To specify whether the server uses TLS for the entire session or only for authentication, see the SSL client option for client-to-server communication, and the **UPDATE SERVER** command with the **SSL** parameter for server-to-server communication.

Beginning in V8.1.2, TLS is used for authentication by default. If you decide to use TLS to encrypt entire sessions, use the protocol only for sessions where it is necessary and add processor resources on the server to manage the increase in network traffic. You can also try other options. For example, some networking devices such as routers and switches 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 information

[Securing communications](#)

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 IBM Storage Protect encryption, or an external device such as tape for encryption.

If you choose IBM Storage Protect to encrypt the data, extra computing resources are required at the client that might affect the performance of backup and restore processes.

Related information

[Data encryption considerations for cloud-container storage pools in IBM Storage Protect](#)

Planning firewall access

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

Table 11 on page 25 describes the ports that are used by the server, client, and Operations Center.

Table 11. Ports that are used by the server, client, and Operations Center			
Item	Default	Direction	Description
Base port (TCP PORT)	1500	Outbound/ inbound	Each server instance requires a unique port. You can specify an alternative port number instead of using the default. The TCP PORT option listens for both TCP/IP and SSL-enabled sessions from the client. For administrative client traffic, you can use the TCP ADMIN PORT and ADMIN NONCLIENT PORT options to set port values.
SSL-only port (SSL TCP PORT)	No default	Outbound/ inbound	This port is used if you want to restrict communication on the port to SSL-enabled sessions only. To support both SSL and non-SSL communications, use the TCP PORT or TCP ADMIN PORT options.
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.

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

Item	Default	Direction	Description
NDMP	No default	Inbound/ outbound	<p>The 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.</p> <p>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.</p> <p>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.</p>
Replication	No default	Outbound/ inbound	<p>The port and protocol for the outbound port for replication are set by the DEFINE SERVER command that is used to set up replication.</p> <p>The inbound ports for replication are the TCP ports and SSL ports that the source replication server names in the DEFINE SERVER command.</p>
Client schedule port	Client port: 1501	Outbound	<p>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.</p>
Long running sessions	KEEPALIVE setting: YES	Outbound	<p>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.</p>
Operations Center	HTTPS: 11090	Inbound	<p>These ports are used for the Operations Center web browser. You can specify an alternative port number.</p>
Client management service port	Client port: 9028	Inbound	<p>The client management service port must be accessible from the Operations Center. Ensure that firewalls cannot prevent connections. The client management service uses the TCP port of the server for the client node for authentication by using an administrative session.</p>

Part 2. Multisite disk implementation of a data protection solution

The multisite disk solution is configured at two sites and uses data deduplication and replication.

Implementation roadmap

The following steps are required to set up a 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 server instance.
 - e. Prepare file systems for IBM Storage Protect.
2. Install the server and Operations Center.
3. Configure the 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 IBM Storage Protect 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.

Tip: Avoid conflicts in managing administrative IDs and client option sets by identifying the IDs and option sets that will be replicated to the target replication server and the IDs and option sets that will be managed in an enterprise configuration. You cannot define an administrative user ID for a registered node if an administrative ID exists for the same node.
 - 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. Configure for SSL communication between the hub and spoke server.
 - b. Add the second server as a spoke.
 - c. Enable replication.
6. Complete the implementation.

Setting up the system

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

Configuring the storage hardware

To configure your storage hardware, review general guidance for disk systems and IBM Storage Protect.

Procedure

1. Provide a connection between the server and the storage devices by following these guidelines:
 - Use a switch or direct connection for Fibre Channel connections.
 - Consider the number of ports that are connected and account for the amount of bandwidth that is needed.
 - Consider the number of ports on the 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 planned properly to ensure optimal performance. See [“Planning for storage” on page 21](#) for more information.
4. Ensure that the server system has access to disk volumes that are created. Complete the following steps:
 - a) If the system is connected to a Fibre Channel switch, zone the server to see the disks.
 - b) Map all of the volumes to tell the disk system that this specific server is allowed to see each disk.

Related information

[Configuring storage](#)

Installing the server operating system

Install the operating system on the server system and ensure that IBM Storage Protect 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 7.2 TL4 or TL5, or later version 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:

```
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
```

Server performance can be affected by the Olson time zone definition.

5. To optimize performance, 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
```

6. In the `.profile` file of the instance user, verify that the following environment variable is set:

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

In later versions of the IBM Storage Protect server, this value is set automatically when the server is started. If the instance user is not available, complete this step later, when the instance user becomes available.

7. Set the system to create full application core files. Issue the following command:

```
chdev -l sys0 -a fullcore=true -P
```

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

- For communications with the server, open port 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 that you are using are open.

9. Enable TCP high-performance enhancements. Issue the following command:

```
no -p -o rfc1323=1
```

10. For optimal throughput and reliability, bond two 10 Gb Ethernet ports together for a medium system and four 10 Gb Ethernet ports for a large system. Use the System Management Interface Tool (SMIT) to bond the ports together by using Etherchannel.

The following settings were used during testing:

mode	8023ad	
auto_recovery	yes	Enable automatic recovery after failover
backup_adapter	NONE	Adapter used when whole channel fails
hash_mode	src_dst_port	Determines how outgoing adapter is chosen
interval	long	Determines interval value for IEEE
		802.3ad mode
mode	8023ad	EtherChannel mode of operation
netaddr	0	Address to ping
no_loss_failover	yes	Enable lossless failover after ping failure
num_retries	3	Times to retry ping before failing
retry_time	1	Wait time (in seconds) between pings
use_alt_addr	no	Enable Alternate EtherChannel Address
use_jumbo_frame	no	Enable Gigabit Ethernet Jumbo Frames

11. Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in [Table 12 on page 30](#). 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 will be installed on the internal hard disks. Configure the internal hard disks by 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 7.8 or later or Version 8.5 or later, according to the manufacturer instructions.

Important: Alternatively, you can also choose to install the following operating systems on the server system:

- SUSE Linux Enterprise Server 15 or later version
- Ubuntu 18.04 LTS or later version

Obtain a bootable DVD or .ISO image that contains Red Hat Enterprise Linux at a supported version and start your system from this media. See the following guidance for installation options. If an item is not mentioned in the following list, leave the default selection.

- a) After you start the operating system installation media, choose **Install or upgrade an existing system** from the menu.
- b) On the Welcome screen, select **Test this media & install Red Hat Enterprise Linux 8.x**.
- c) Select your language and keyboard preferences.
- d) Select your location to set the correct timezone.
- e) Select **Software Selection** and then on the next screen, select **Server with GUI**.
- f) From the installation summary page, click **Installation Destination** and verify the following items:
 - 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 is completed, 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 produced the following result:

```
[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. Bond two ports for a medium system and four ports for a large system. This can be accomplished by creating a Link Aggregation Control Protocol (LACP) network connection, which aggregates several subordinate ports into a single logical connection. The preferred method is to use a bond mode of 802.3ad, **miimon** setting of 100, and a **xmit_hash_policy** setting of layer3+4.

Restriction: To use an LACP network connection, you must have a network switch that supports LACP.

For additional instructions about configuring bonded network connections with Red Hat Enterprise Linux Version 7, see [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 the 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 example:

```
/dev/sr0 on /mnt type iso9660
```

- c) Change to the YUM repository directory by issuing the following command:

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

For RHEL 8:

```
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 rhel78_dvd.repo
```

- g) Add the following lines to the new repo file. The **baseurl** parameter specifies your directory mount point:

```
[rhel78_dvd]
name=DVD Redhat Enterprise Linux 7.8
baseurl=file:///mnt
enabled=1
gpgcheck=0
```

For RHEL 8:

```
[InstallMedia-BaseOS]
name=Red Hat Enterprise Linux 8.2.0
mediaid=None
metadata_expire=-1
gpgcheck=0
cost=500
enabled=1
baseurl=file:///mnt/BaseOS/

[InstallMedia-AppStream]
name=Red Hat Enterprise Linux 8.2.0
mediaid=None
metadata_expire=-1
gpgcheck=0
cost=500
enabled=1
baseurl=file:///mnt/AppStream/
```

- h) Install additional prerequisite software packages, by issuing the **yum** command.
For example:

```
yum install ksh.x86_64
yum install sysstat
For RHEL 8:
yum install libnsl
```

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:

```
umount /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 rhel78_dvd.repo rhel78_dvd.repo.orig
```

- d) Rename the original file to the original name:

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

6. Determine whether 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 on page 33](#) to determine whether 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.

Tip: Automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings. On Linux, the Db2 database software automatically adjusts interprocess communication (IPC) kernel parameter values to the preferred settings. For more information about kernel parameter settings, search for Linux kernel parameters in the [Version 11.5 product documentation](#).

Table 13. Linux kernel parameter optimum settings	
Parameter	Description
kernel.shmmni	The maximum number of segments.
kernel.shmmax	The maximum size of a shared memory segment (bytes). This parameter must be set before automatically starting the IBM Storage Protect server on system startup.
kernel.shmall	The maximum allocation of shared memory pages (pages).
kernel.sem There are four values for the kernel.sem parameter.	(SEMMSL) The maximum semaphores per array.
	(SEMMNS) The maximum semaphores per system.
	(SEMOPM) The maximum operations per semaphore call.
	(SEMMNI) The maximum number of arrays.
kernel.msgmni	The maximum number of system-wide message queues.
kernel.msgmax	The maximum size of messages (bytes).
kernel.msgmnb	The default maximum size of queue (bytes).
kernel.randomize_va_space	The kernel.randomize_va_space parameter configures the use of memory ASLR for the kernel. Enable ASLR for V7.1 and later servers. To learn more details about the Linux ASLR and Db2, see technote 1365583 .
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® product information .
vm.overcommit_memory	The vm.overcommit_memory parameter influences how much virtual memory the kernel permits allocating. For more information about kernel parameters, see the Db2 product information .

7. Open firewall ports to communicate with the server. 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 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 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 process resource limits, also known as *ulimits*, are set according to guidelines in [Table 14 on page 34](#). 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 2016 or 2019 Standard Edition on the server system and prepare the system for installation and configuration of the IBM Storage Protect server.

Procedure

1. Install Windows Server 2016 or 2019 Standard Edition, according to the manufacturer instructions.
2. Change the Windows account control policies by completing the following steps.
 - a) Open the Local Security Policy editor by running `secpol.msc`.
 - b) Click **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 Server updates.
 - b) If required, update the FC and Ethernet HBA device drivers to newer levels.
5. Open the default TCP/IP port, 1500, for communications with the IBM Storage Protect server.
For example, issue the following command:

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

Configuring multipath I/O

You can 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.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
```

The output is similar to the following example:

```
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, a device ID could be **hdisk4**. Save the list of device IDs to use when you create file systems for the IBM Storage Protect 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
```

On an IBM Storwize system, the following information is an example of what is shown for each device:

```
hdisk4 00f8cf083fd97327 None active
3321360050763008101057800000000000003004214503IBMfcp
```

In the example, **60050763008101057800000000000030** is the UID for the volume, as reported by the Storwize management interface.

To verify disk size in megabytes and compare the value 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.
If the `multipath.conf` file does not exist, you can create it by issuing the following command:

```
mpathconf --enable
```

The following parameters were set in `multipath.conf` for testing on an IBM FlashSystem® storage system:

```
defaults {
    user_friendly_names no
}

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

```

        path_selector "round-robin 0"
        prio "alua"
        path_checker "tur"
        failback "immediate"
        no_path_retry 5
        rr_weight uniform
        rr_min_io_rq "1"
        dev_loss_tmo 120
    }
}

```

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, 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
36005076802810c5098000000000000012 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 36005076802810c5098000000000000012.

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. For IBM FlashSystem devices, use the Microsoft Device Specific Module (MSDSM). For installation instructions, see the IBM FlashSystem documentation https://www.ibm.com/support/knowledgecenter/STHGUJ_8.3.1/com.ibm.storwize.v5000.831.doc/svc_w2kmpio_21oxvp.html
2. To verify that disks are visible to the operating system and are managed by multipath I/O, open a Microsoft Windows Power Shell command prompt and issue the following command:

```

mpclaim -e

```

3. Review the mpclaim output and ensure that the IBM storage is reported as under MPIO control.

"Target H/W Identifier	"	Bus Type	MPIO-ed	ALUA Support
"IBM 2145	"	SAS	YES	Implicit Only

4. Additional details of attach disk devices can be obtained using the Windows wmic command.

```
wmic diskdrive get
```

5. To bring new disks online 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 and directories for the server instance

Create the user ID that owns the IBM Storage Protect server instance and create the directories that the server instance needs for database and recovery logs.

Before you begin

Review the information about planning space for the server before you complete this task. See [Worksheets for planning details for the server](#).

Procedure

1. Create the user ID that will own the server instance.

You use this user ID when you create the server instance in a later step.

Linux | **AIX**

Create a user ID and group that will be the owner of the server instance.

- a. The following commands can be run from an administrative user ID that will set up the user and group. Create the user ID and group in the home directory of the user.

Restriction: In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. 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.

For example, create user ID *tsminst1* in group *tsmsrvrs*. The following examples show how to create this user ID and group using operating system commands.

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



```
Linux groupadd tsmisvrs -g 1111
useradd -d /home/tsminst1 -u 2222 -g 1111 -s /bin/bash tsminst1
passwd tsminst1
```

Restriction: IBM Db2 does not support direct operating system user authentication through LDAP.

- b. Log off, 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 that will be the owner of the IBM Storage Protect server instance. A user ID can own more than one IBM Storage Protect server instance. Identify the user account that will own the server instance.

When the server is started as a Windows service, this account is the one that the service will log on to. The user account must have administrative authority on the system. One user account can own more than one server instance.

If you have multiple servers on one system and want to run each server with a different user account, create a new user account in this step.

Create the user ID.

Restriction: The user ID must comply with the following rule:

In the user ID, only lowercase letters (a-z), numerals (0-9), and the underscore character (_) can be used. The user ID must be 30 characters or less, and 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.

- a. Use the following operating system command to create the user ID:

```
net user user_ID * /add
```

You are prompted to create and verify a password for the new user ID.

- b. Issue the following operating system commands to add the new user ID to the Administrators groups:

```
net localgroup Administrators user_ID /add
net localgroup DB2ADMNS user_ID /add
```

2. Create directories that the server requires.

Linux

AIX

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	<code>mkdir /tsminst1</code>	

Create empty directories for each item in the table and ensure that the directories are owned by the new user ID you just created. Mount the associated storage to each directory for the active log, archive log, and database directories. (*continued*)

Item	Example commands for creating the directories	Your directories
The database directories	mkdir /tsminst1/ TSMdbspace00 mkdir /tsminst1/ TSMdbspace01 mkdir /tsminst1/ TSMdbspace02 mkdir /tsminst1/ TSMdbspace03	
Active log directory	mkdir /tsminst1/TSMalog	
Archive log directory	mkdir /tsminst1/TSMarchlog	
Optional: Directory for the log mirror for the active log	mkdir /tsminst1/ TSMlogmirror	
Optional: Secondary archive log directory (failover location for archive log)	mkdir /tsminst1/ TSMarchlogfailover	

Windows

Create empty directories for each item in the table and ensure that the new user ID you just created has read/write permission to the directories. The database, archive log, and active log must reside on different physical volumes.

Item	Example commands for creating the directories	Your directories
The <i>instance directory</i> for the server, which is a directory that will contain files specifically for this server instance (the server options file and other server-specific files)	mkdir C:\tsminst1	
The database directories	mkdir C:\tsminst1\TSMdbspace00 mkdir C:\tsminst1\TSMdbspace01 mkdir C:\tsminst1\TSMdbspace02 mkdir C:\tsminst1\TSMdbspace03	
Active log directory	mkdir \tsminst1\TSMalog	
Archive log directory	mkdir C:\tsminst1\TSMarchlog	
Optional: Directory for the log mirror for the active log	mkdir C:\tsminst1\TSMlogmirror	

Create empty directories for each item in the table and ensure that the new user ID you just created has read/write permission to the directories. The database, archive log, and active log must reside on different physical volumes. (<i>continued</i>)		
Item	Example commands for creating the directories	Your directories
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir C:\tsminst1\TSMarchlogfailover</code>	

When a server is initially created by using the **DSMSERV FORMAT** utility or the configuration wizard, a server database and recovery log are created. In addition, files are created to hold database information that is used by the database manager.

3. Log off the new user ID.

Preparing file systems for the server

You must complete file system configuration for the disk storage to be used by the server.

AIX systems

You must create volume groups, logical volumes, and file systems for the 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. 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 IBM Storage Protect 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.

System size: The following commands are based on the medium system configuration. For small and large systems, you must adjust the syntax as required.

```
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 **mk1v** 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 IBM Storage Protect database on a medium system, issue the following commands:

```
mk1v -y tsmdb00 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk2
mk1v -y tsmdb01 -t jfs2 -u 1 -x 1631 tsmdb 1631 hdisk3
mk1v -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 following example of command output shows that the amount of used space is typically 1%:

```
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 and directories for the server instance”](#) on [page 38](#) has read and write access to the directories for the IBM Storage Protect server.

Linux systems

You must format ext4 or xfs file systems on each of the disk LUNs to be used by the IBM Storage Protect server.

Procedure

1. Using the list of device IDs that you generated previously, issue the **mkfs** command to create and format a file system for each storage LUN device. Specify the device ID in the command. See the following examples.

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/360050763008101057800000000000002c3
```

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 file systems.

Issue the **mkdir** command for each 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/36005076802810c50980000000000012 /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 following example on an IBM Storwize system shows that the amount of used space is typically 1%:

```
[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 and directories for the server instance”](#) on page 38 has read and write access to the directories for IBM Storage Protect.

Windows systems

You must format New Technology File System (NTFS) file systems on each of the disk LUNs to be used by the IBM Storage Protect server.

Procedure

1. Create mount point directories for file systems.

Issue the **md** command for each 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.

Installing the server and Operations Center

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

Installing on AIX and Linux systems

Install the IBM Storage Protect 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 45 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.

For space requirements, see the download document at [technote 7169125](#) and [technote 7169558](#).

3. Go to [Passport Advantage®](#) and 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 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.

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

- After you install the server and before you customize it for your use, go to the [support site](#). Click **Support and downloads** and apply any applicable fixes.

Related information

[Installing the server components](#)

AIX Installing prerequisite RPM files for the graphical wizard

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](#) 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 IBM Storage Protect 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.
For space requirements, see the download document at [technote 7168813](#).
2. Go to [Passport Advantage](#) and 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.

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

- After you install the server and before you customize it for your use, go to the [support site](#). Click **Support and downloads** and apply any applicable fixes.

Related information

[Installing the server components](#)

Configuring the server and the Operations Center

After you install the components, complete the configuration for the IBM Storage Protect server and the Operations Center.

Configuring the server instance

Use the IBM Storage Protect server instance configuration wizard to complete the initial configuration of the server.

Before you begin

Ensure that the following requirements are met:

Linux | **AIX**

- The system where you installed IBM Storage Protect 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 IBM Storage Protect 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 changed any settings in the preceding steps, 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 that 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**.
 - d. Select **New Rule**.
 - 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 changed any settings in the preceding steps, 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.
 - **Linux | AIX** 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 > IBM Storage Protect > Configuration Wizard**.
2. Follow the instructions to complete the configuration.

Use the information that you recorded in “[Planning worksheets](#)” on [page 10](#) during IBM Storage Protect system set up to specify directories and options in the wizard.

 - **Linux | AIX** On the **Server Information** window, set the server to start automatically by using the instance user ID when the system boots.
 - **Windows** By using the configuration wizard, the server is set to start automatically when rebooted.

Installing the backup-archive client

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

Procedure

- To install the backup-archive client, follow the installation instructions for your operating system.
 - [Install UNIX and Linux backup-archive clients](#)

- Installing the Windows client for the first time

Setting options for the server

Review the server options file that is installed with the IBM Storage Protect 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 values in the following table 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	TCP/IP	TCP/IP	TCP/IP
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
NUMOPENVOLSAALLOWED	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. To make updates, 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 any server 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.
 - On each line, enter only one option and the specified value for the option.
 - If an option occurs in multiple entries 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 information

[Server options reference](#)

[SETOPT \(Set a server option for dynamic update\)](#)

Configuring secure communications with Transport Layer Security

To encrypt data and secure communications in your environment, Secure Sockets Layer (SSL) or Transport Layer Security (TLS) is enabled on the IBM Storage Protect server and backup-archive client. An SSL certificate is used to verify communication requests between the server and client.

About this task

Beginning with IBM Storage Protect Version 8.1.2, SSL is enabled by default, and the IBM Storage Protect server and backup-archive client are automatically configured to communicate with each other by using protocol TLS Version 1.2 or later.

As shown in the following figure, you can manually configure secure communications between the 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. Alternatively, you can obtain and transfer a unique certificate that is signed by a certificate authority (CA).



For more information about configuring the server and clients for SSL or TLS communications, see [Configuring storage agents, servers, clients, and the Operations Center to connect to the server by using SSL](#).

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 server that you want to designate as a hub server
- Login credentials for an administrator ID that is defined for that server

Procedure

1. 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 by using self-signed certificates](#).

2. 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.

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 email reports”](#) on page 84.

Securing communications between the Operations Center and the hub server

To secure communications between the Operations Center and the hub server, add the Transport Layer Security (TLS) 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 certificates that the Operations Center can access. During the installation of the Operations Center, you must create a password for the truststore file. To secure communications between the Operations Center and the hub server, you must use the same password to add the certificate of the hub server to the truststore file. If you do not remember this password, you must now re-create and configure the truststore file. For instructions, see "Deleting and reassigning the password for the Operations Center truststore file" in IBM Documentation.

The following figure illustrates the components for setting up a Secure Sockets Layer (SSL) connection between the hub server and the Operations Center.



About this task

This procedure provides steps to implement secure communications by using self-signed certificates.

If you use certificates that are signed by a certificate authority (CA), follow the instructions in "Securing communications between the Operations Center and the hub server by using CA-signed certificates" in IBM Documentation.

Procedure

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

1. Stop the Operations Center web server.
2. Open the operating system command line on the system where the Operations Center is installed, and change to the following directory:

- **Linux | AIX** `installation_dir/ui/jre/bin`
- **Windows** `installation_dir\ui\jre\bin`

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

3. Open the IBM Key Management window by issuing the following command:

```
ikeyman
```

4. Click **Key Database File > Open**.
5. Click **Browse**, and go to the following directory, where *installation_dir* represents the directory in which the Operations Center is installed:

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

6. In the `guiServer` directory, select the `gui-truststore.jks` file.
7. Click **Open**, and click **OK**.
8. Enter the password for the truststore file, and click **OK**.
9. In the **Key database content** area of the IBM Key Management window, click the arrow, and select **Signer Certificates** from the list. Click **Add**.
10. In the Open window, click **Browse**, and go to the hub server instance directory, which was specified by the administrator who created the instance. For example:

- **Linux | AIX** `home/tsminst1`
- **Windows** `c:\Program Files\Tivoli\TSM\server1`

The directory contains the `cert256.arm` certificate.

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's instance directory to the following directory on the computer where the Operations Center is installed:

- **Linux | AIX** `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.

11. Select the `cert256.arm` certificate as the SSL certificate.
12. Click **Open**, and click **OK**.
13. Enter a label for the certificate. For example, enter the name of the hub server.
14. 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.
15. Close the IBM Key Management window.
16. Start the Operations Center web server.

When you connect to the Operations Center for the first time, you are prompted to identify the IP address or network name of the hub server, and the port number for communicating with the hub server. If the `ADMINONCLIENTPORT` server option is enabled for the IBM Storage Protect server, enter the port number that is specified by the `TCPADMINPORT` server option. If the `ADMINONCLIENTPORT` server option is not enabled, enter the port number that is specified by the `TCPPORT` server option.

Related tasks

[Starting and stopping the web server](#)

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

Registering the product license


To register your license for the IBM Storage Protect product, 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 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 settings icon  and clicking **Command Builder**.
3. Issue the **REGISTER LICENSE** command.
For example, to register a base IBM Storage Protect 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. The server stores license information in the NODELOCK file, which is in the directory from which the server is started.
-  If you change the processor chip that is associated with the server on which the server is installed.

Related information

[REGISTER LICENSE \(Register a new license\)](#)

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 [“Planning worksheets” on page 10](#) 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 a 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](#)

Before you add a client, ensure that appropriate rules are specified for backup and archive operations for 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.

Defining schedules for server maintenance activities

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

About this task

Schedule server maintenance operations to run after client backup operations. You can control the timing of schedules 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.

Operation	Schedule
Processing for database and disaster recovery files	<ul style="list-style-type: none"> 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

After you configure the device class for the database backup operations, create schedules for database backup and other required maintenance operations by using the **DEFINE SCHEDULE** command. Depending on the size of your environment, you might need to adjust the start times for each schedule in the example.

1. Define a device class for the backup operations.

For example, use the **DEFINE DEVCLASS** command to create a device class that is named **DBBACK_FILEDEV**:

```
define devclass dback_filedev devtype=file
directory=db_backup_directories
```

where *db_backup_directories* is a list of the directories that you created for the database backup.

Linux | **AIX** For example, if you have four directories for database backups, starting with /tsminst1/TSMbkup00, issue the following command:

```
define devclass dback_filedev devtype=file
directory=/tsminst1/TSMbkup00,
/tsminst1/TSMbkup01,/tsminst1/TSMbkup02,
/tsminst1/TSMbkup03"
```

Windows For example, if you have four directories for database backups, starting with C:\tsminst1\TSMbkup00, issue the following command:

```
define devclass dback_filedev devtype=file
directory="c:\tsminst1\TSMbkup00,
c:\tsminst1\TSMbkup01,c:\tsminst1\TSMbkup02,
c:\tsminst1\TSMbkup03"
```

2. Set the device class for automatic database backup operations. Use the **SET DBRECOVERY** command to specify the device class that you created in the preceding step.

For example, if the device class is **dback_filedev**, issue the following command:

```
set dbrecovery dback_filedev
```

3. Create schedules for the maintenance operations by using the **DEFINE SCHEDULE** command. See the following table for the required operations with examples of the commands.

Tip: You create the schedule for replication separately in a later step, when you use the Operations Center to configure replication.

Operation	Example command
Back up the database.	<p>Create a schedule to run the BACKUP DB command. If you are configuring a small system, set the COMPRESS parameter to YES.</p> <p>For example, on a small system, issue the following command to create a backup schedule that uses the new device class:</p> <pre>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</pre>
Back up the device configuration information.	<p>Create a schedule to run the BACKUP DEVCONFIG command:</p> <pre>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</pre>
Back up the volume history.	<p>Create a schedule to run the BACKUP VOLHISTORY command:</p> <pre>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</pre>
Remove older versions of database backups that are no longer required.	<p>Create a schedule to run the DELETE VOLHISTORY command:</p> <pre>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</pre>
Remove objects that exceed their allowed retention.	<p>Create a schedule to run the EXPIRE INVENTORY command.</p> <p>Set the RESOURCE parameter based on the system size that you are configuring:</p> <ul style="list-style-type: none"> • Small systems: 10 • Medium systems: 30 • Large systems: 40 <p>For example, on a medium-sized system, issue the following command to create a schedule that is named EXPINVENTORY:</p> <pre>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</pre>

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 information

[DEFINE SCHEDULE \(Define a schedule for an administrative command\)](#)

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 53.

Installing and configuring backup-archive clients

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

Procedure

- To install the backup-archive client, follow the installation instructions for your operating system.
 - [Install UNIX and Linux backup-archive clients](#)
 - [Installing the Windows client for the first time](#)

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.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see [technote 7048963](#).

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see [Authenticating users by using an Active Directory database](#).

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the **REGISTER NODE** command and specify the **USERID** parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see [Register a node](#).

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:

- a. On the Operations Center menu bar, click **Clients**.
- b. In the Clients table, click **+Client**.
- c. Complete the steps in the **Add Client** wizard:
 - i) Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the **Enable** check box.
 - ii) 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.
 - iii) Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
 - iv) Set how risks are displayed for the client by specifying the at-risk setting.
 - v) 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 backup-archive clients and makes the information available to the Operations Center for basic monitoring capability.

Before you begin

Restriction: The following procedure is applicable only if you installed IBM Storage Protect client management service as a separate component earlier than IBM Storage Protect 8.1.13. In IBM Storage Protect 8.1.13, the installation of a separate client management service package was deprecated and the feature that the client management service provided was integrated into the backup-archive client package.

For information about collecting diagnostic information for the backup-archive client version 8.1.13 and later, see *Collecting diagnostic information when the Operations Center and backup-archive clients are at 8.1.13 or later* in IBM Documentation.

- Review the *Requirements and limitations for IBM Storage Protect client management services* in IBM Documentation.
- Before you install the client management service, ensure that a successful connection was established between the backup-archive client and the server. The server truststore file that the client uses does not have the server Secure Sockets Layer (SSL) certificate until the client system has connected to the server.

Procedure

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

1. 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>-IBM Storage Protect-CMS-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 Storage Protect Client Management Services.

Related information

[Configuring the client management service for custom client installations](#)

Verifying that 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. The following procedure is applicable only if you installed IBM Storage Protect client management service as a separate component.

Procedure

On the client system, at the command line, run the following commands to view the configuration of the client management service:

- On Linux client systems, issue the following command:

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

where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
/opt/tivoli/tsm/cms/bin/CmsConfig.sh list
```

The output is similar to the following text:

```
Listing CMS configuration
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: /opt/tivoli/tsm/client/ba/bin/dsm.sys
  Log File: /opt/tivoli/tsm/client/ba/bin/dsmerror.log
             en_US MM/dd/yyyy HH:mm:ss Windows-1252
  Log File: /opt/tivoli/tsm/client/ba/bin/dsmsched.log
             en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

- On Windows client systems, issue the following command:

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

where *client_install_dir* is the directory where the backup-archive client is installed. For example, with the default client installation, issue the following command:

```
C:\Program Files\Tivoli\TSM\cms\bin\CmsConfig.bat list
```

The output is similar to the following text:

```
Listing CMS configuration
server1.example.com:1500 NO_SSL HOSTNAME
Capabilities: [LOG_QUERY]
  Opt Path: C:\Program Files\Tivoli\TSM\baclient\dsm.opt
  Log File: C:\Program Files\Tivoli\TSM\baclient\dsmerror.log
             en_US MM/dd/yyyy HH:mm:ss Windows-1252
  Log File: C:\Program Files\Tivoli\TSM\baclient\dsmsched.log
             en_US MM/dd/yyyy HH:mm:ss Windows-1252
```

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:

- On Linux client systems:

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

- On Windows client systems:

```
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 to configure this file, see [Configuring the client management service for custom client installations](#). You can use the **CmsConfig verify** command to verify that a node definition is correctly created in the `client-configuration.xml` file.

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

Restriction: The following procedure is applicable only if you installed IBM Storage Protect client management service as a separate component.

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 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	<code>https://server.example.com</code>
With DNS host name and non-default port	<code>https://server.example.com:1599</code>
With IP address and non-default port	<code>https://192.0.2.0:1599</code>

4. Click **Save**.

What to do next

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

Configuring the second server

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

Procedure

Complete the instructions in the following sections:

1. Configure a second server that is the same as the first server by completing the instructions in the following sections:

- a) [“Setting up the system” on page 28](#)
- b) [“Installing the server and Operations Center” on page 44](#)

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.

- c) [“Configuring the server and the Operations Center” on page 46](#)

Skip the tasks for configuring the Operations Center.

- d) [“Installing and configuring backup-archive clients” on page 56](#)
2. [“Configuring SSL communications between the hub server and a spoke server” on page 60](#)
 3. [“Adding the second server as a spoke” on page 62](#)
 4. [“Enabling replication” on page 62](#)

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 Transport Layer Security (TLS) protocol, you must define the certificate of the spoke server to the hub server.

About this task

The hub server receives status and alert information from the spoke server and shows this information in the Operations Center. To receive the status and alert information from the spoke server, the certificate of the spoke server must be added to the truststore file of the hub server. You must also configure the Operations Center to monitor the spoke server.

To enable other functions of the Operations Center, such as the automatic deployment of client updates, the certificate of the hub server must be added to the truststore file of the spoke server.

Procedure

1. Complete the following steps to define the certificate of the spoke server to the hub server:
 - a) On the spoke server, change to the directory of the spoke server instance.
 - b) Verify the certificates in the key database file of the spoke server. Issue the following command:

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

- c) Securely transfer the cert256.arm file of the spoke server to the hub server.
- d) On the hub server, change to the directory of the hub server instance.

- e) Define the spoke server 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 certificate:

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

2. Complete the following steps to define the certificate of the hub server to the spoke server:
- a) On the hub server, change to the directory of the hub server instance.
- b) Verify the certificates in the key database file of the spoke server. Issue the following command:

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

- c) Securely transfer the *cert256.arm* file of the hub server to the spoke server.
- d) On the spoke server, change to the directory of the spoke server instance.
- e) Define the hub server certificate to the spoke server. Issue the following command from the spoke server instance directory, where *hub_servername* is the name of the hub server, and *hub_cert256.arm* is the file name of the hub server certificate:

```
gsk8capicmd_64 -cert -add -db cert.kdb -stashed -format ascii -trust enable  
-label hub_servername -file hub_cert256.arm
```

3. Restart the hub server and the spoke server.
4. Complete the following steps to define the spoke server to the hub server, and the hub server to the spoke server.
- a) Issue the following commands on both the hub server and the spoke server:

```
SET SERVERPASSWORD server_password  
SET SERVERHLADDRESS ip_address  
SET SERVERLLADDRESS tcp_port
```

- b) On 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
```

- c) On the spoke server, issue the **DEFINE SERVER** command, according to the following example:

```
DEFINE SERVER hub_servername HLA=hub_address  
LLA=hub_SSLTCPADMINPort SERVERPA=hub_serverpassword
```

Tip: By default, server communication is encrypted except when the server is sending or receiving object data. Object data is sent and received by using TCP/IP. By choosing not to encrypt the object data, server performance is similar to communication over a TCP/IP session and the session is secure. To encrypt all communication with the specified server, even when the server is sending and receiving object data, specify the **SSL=YES** parameter on the **DEFINE SERVER** command.

5. Complete the following steps to configure the Operations Center to monitor the spoke server:
- a) On the Operations Center menu bar, click **Servers**.
- The spoke server has a status of "Unmonitored." This status means that, although this server was defined to the hub server by using the **DEFINE SERVER** command, the server is not yet configured as a spoke.
- b) Click the spoke server to highlight the item, and click **Monitor Spoke**.

Related information

[DEFINE SERVER \(Define a server for server-to-server communications\)](#)

[QUERY OPTION \(Query server options\)](#)

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 replication 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 replication server as the first server that you configured for the multisite disk solution. The target replication 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 53.
 - 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.

Completing the implementation

After the IBM Storage Protect 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 for which you want to back up the database. Click **Back Up**, and follow the instructions in the **Back Up Database** window.
 - c) Verify that the backup operations completed successfully with no warning or error messages.

Tip: Alternatively, you can use the backup-archive client GUI to back up client data and you can backup the server database by issuing **BACKUP DB** command from an administrative command-line.

2. Set up monitoring for your solution by following the instructions in [Part 3, “Monitoring a multisite disk solution,” on page 65.](#)

Part 3. Monitoring a multisite disk solution

After you implement a multisite disk solution, monitor the solution to ensure correct operation.

About this task

After you implement a multisite disk solution with IBM Storage Protect, monitor the solution daily and periodically to identify existing and potential issues. The information that you gather can be used to troubleshoot problems and optimize system performance.

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.

Tip: If you plan to diagnose issues with backup-archive clients on Linux or Windows operating systems, install IBM Storage Protect 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 57](#).

Procedure

1. Complete daily monitoring tasks. For instructions, see [“Daily monitoring checklist” on page 65](#).
2. Complete periodic monitoring tasks. For instructions, see [“Periodic monitoring checklist” on page 77](#).
3. To verify that your IBM Storage Protect solution complies with licensing requirements, follow the instructions in [“Verifying license compliance” on page 82](#).
4. To set up Operations Center to generate email status reports, see [“Tracking system status by using email reports” on page 84](#)

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 for a multisite disk solution,” on page 87](#). The following resources are also available:


- To resolve performance issues, see [Performance](#).
- To resolve other types of issues, see [Troubleshooting](#).


Daily monitoring checklist

To ensure that you are completing the daily monitoring tasks for your IBM Storage Protect 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 settings 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
<p>Watch for security notifications, which can indicate a ransomware attack.</p>	<p>If a potential ransomware attack is detected in the IBM Storage Protect environment, a security notification message is displayed in the foreground of the Operations Center. For more information, click the message to open the Security Notifications page.</p>	<p>On the Security Notifications page, you can take the following actions:</p> <ul style="list-style-type: none"> • View notification details by client. <p>Restriction: Notifications are available only for backup-archive clients and IBM Storage Protect for Virtual Environments clients.</p> <ul style="list-style-type: none"> • Acknowledge a security notification by selecting it and clicking Acknowledge. When you acknowledge a security notification, a check mark is added to the Acknowledged column of the Security Notifications page for the selected client. The standard by which a notification is acknowledged is determined by your organization. A check mark might mean that you investigated the issue and determined that it is a false positive. Or it might mean that a problem exists and is being resolved. • Assign a security notification to an administrator by selecting the security notification and clicking Assign. To view the assignment, the administrator must sign in to the Operations Center and click Overviews > Security. If you are not certain whether the administrator regularly monitors the Security Notifications page, notify the administrator about the assignment. • If the notification is a false positive, you can select the security notification and click Reset. The security notification is deleted. Historical data that is used for baseline comparisons with the most recent backup operation is deleted. A new baseline is calculated going forward. • Optionally, you can disable security notifications by using the SET SECURITYNOTIF command.

Table 15. Daily monitoring tasks (continued)


Task	Basic procedures	Advanced procedures and troubleshooting information
<p>1 Determine whether clients are at risk of being unprotected due to failed or missed backup operations.</p>	<p>To verify whether clients are at risk, in the Clients area, look for an At risk notification. To view details, click the Clients area.</p> <p> Attention: If the At risk percentage is much greater than usual, it might indicate a ransomware attack. A ransomware attack can cause backup operations to fail, thus placing clients at risk. For example, if the percentage of clients at risk is normally between 5% and 10%, but the percentage increases to 40% or 50%, investigate the cause.</p> <p>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:</p> <ol style="list-style-type: none"> 1. In the Clients table, select the client and click Details. 2. To diagnose an issue, click Diagnosis. 	<p>For clients that do not have the client management service installed, access the client system to review the client error logs.</p>
<p>2 Determine whether client-related or server-related errors require attention.</p>	<p>To determine the severity of any reported alerts, in the Alerts area, hover over the columns.</p>	<p>To view additional information about alerts, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Alerts area. 2. In the Alerts table, select an alert. 3. In the Activity Log pane, review the messages. The pane displays related messages that were issued before and after the selected alert occurred.
<p>3 Determine whether servers that are managed by the Operations Center are available to provide data protection services to clients.</p>	<ol style="list-style-type: none"> 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. 	<p>Tip: If you detect an issue that is related to server properties, update the server properties:</p> <ol style="list-style-type: none"> 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
<p>4 Determine whether sufficient space is available for the server inventory, which consists of the server database, active log, and archive log.</p>	<ol style="list-style-type: none"> Click the Servers area. In the Status column of the table, view the status of the server and resolve any issues: <ul style="list-style-type: none"> 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. 	<p>You can also look for related alerts on the Alerts page. For additional instructions about troubleshooting, see Resolving server problems.</p>

Table 15. Daily monitoring tasks (continued)


Task	Basic procedures	Advanced procedures and troubleshooting information
<p>5 Verify server database backup operations.</p>	<p>To determine when a server was most recently backed up, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click the Servers area. 2. In the Servers table, review the Last Database Backup column. 	<p>To obtain more detailed information about backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 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. <p>If a database was not backed up recently (for example, in the last 24 hours), you can start a backup operation:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click the Servers area. 2. In the table, select a server and click Back Up. <p>To determine whether the server database is configured for automatic backup operations, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the menu bar, hover over the settings icon  and click Command Builder. 2. Issue the QUERY DB command: <pre>query db f=d</pre> <ol style="list-style-type: none"> 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.
<p>6 Monitor other server maintenance tasks. Server maintenance tasks can include running administrative command schedules, maintenance scripts, and related commands.</p>	<p>To search for information about processes that failed because of server issues, complete the following steps:</p> <ol style="list-style-type: none"> 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 checkbox that is associated with the process. 	<p>For more information about monitoring processes and resolving issues, see the Operations Center online help.</p>

Table 15. Daily monitoring tasks (continued)


Task	Basic procedures	Advanced procedures and troubleshooting information
<p>7 Verify that the amount of data that was recently sent to and from servers is within the expected range.</p>	<ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • 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. <p> Attention: If the amount of backed-up data is significantly larger than usual, it might indicate a ransomware attack. When ransomware encrypts data, the system perceives the data as being changed, and the changed data is backed up. Thus, backup volumes become larger. To determine which clients are affected, click the Applications, Virtual Machines, or Systems tab.</p> <ul style="list-style-type: none"> • 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
<p>8 Verify that storage pools are available to back up client data.</p>	<ol style="list-style-type: none"> If problems are indicated in the Storage & Data Availability area, click Pools to view the details: <ul style="list-style-type: none"> If the Critical  status is displayed, insufficient space is available in the storage pool, or its access status is unavailable. <div data-bbox="414 541 479 598">  </div> Attention: If the status is critical, investigate the cause: <ul style="list-style-type: none"> If the data deduplication rate for a storage pool drops significantly, it might indicate a ransomware attack. During a ransomware attack, data is encrypted and cannot be deduplicated. To verify the data deduplication rate, in the Storage Pools table, review the value in the % Savings column. If a storage pool unexpectedly becomes 100% utilized, it might indicate a ransomware attack. To verify the utilization, review the value in the Capacity Used column. Hover over the values to see the percentages of used and free space. If the Warning  status is displayed, the storage pool is running out of space, or its access status is read-only. To view the used, free, and total space for your selected storage pool, hover over the entries in the Capacity Used column. 	<p>To view the storage-pool capacity that was used over the past two weeks, select a row in the Storage Pools table and click Details.</p>

Table 15. Daily monitoring tasks (continued)



Task	Basic procedures	Advanced procedures and troubleshooting information
<p>9 Verify that storage devices are available for backup operations.</p>	<p>In the Storage & Data Availability area, in the Volumes section, under the capacity bars, review the status that is reported next to Devices. If a Critical  or Warning  status is displayed for any device, investigate the issue. To view details, click Devices.</p>	<p>Disk devices might have a critical or warning status for the following reasons:</p> <ul style="list-style-type: none"> • For DISK device classes, volumes might be offline or have a read-only access status. The Disk Storage column of the Disk Devices table shows the state of volumes. • For FILE device classes that are not shared, directories might be offline. Also, insufficient free space might be available for allocating scratch volumes. The Disk Storage column of the Disk Devices table shows the state of directories. • For FILE device classes that are shared, drives might be unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. Other columns of the Disk Devices table show the state of the drives and paths.

Table 15. Daily monitoring tasks (continued)




Task	Basic procedures	Advanced procedures and troubleshooting information
<p>10 Monitor node replication processes.</p>	<ol style="list-style-type: none"> 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. <p> Attention: If you notice an unexpected increase in the number of replication failures, it might indicate a ransomware attack. Investigate the cause of the failures.</p> <ol style="list-style-type: none"> 3. To view the amount of data that was replicated over the last two 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. <p> Attention: If you see a drastic, unexpected increase in the replication workload, it might indicate a ransomware attack. Investigate the cause of the increased workload.</p>	<p>For advanced monitoring, view information about running and ended node replication processes by using commands:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. Issue the QUERY REPLICATION command. For instructions, see QUERY REPLICATION (Query node replication processes). If the replication operation was completed successfully, the Total Files To Replicate value matches the Total Files Replicated value. <p>To display messages that are related to a node replication process on a source or target replication server, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, click Servers. 2. Select the source or target replication server and click Details: <ul style="list-style-type: none"> • To view active tasks, click Active Tasks, select the task, and verify that the Running status 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.

Table 15. Daily monitoring tasks (continued)






Task	Basic procedures	Advanced procedures and troubleshooting information
<p>11 Monitor retention sets.</p>	<p>To obtain the overall status of retention sets, view the Retention Sets area on the Operations Center Overview page:</p> <ul style="list-style-type: none"> • The Completed field specifies the number of retention sets that were created in the server database and are tracked in the server inventory. • The Expired field specifies the number of retention sets whose data is expired. • The Deleted field specifies the number of retention sets that were deleted. <p>To view or modify retention rules, click Services > Retention Rules.</p>	<p>For more information about retention sets, click the Retention Sets area to open the Retention Sets page. To view or modify retention set properties, double-click a retention set.</p> <p>For more detailed information, you can run related commands:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To determine which retention set creation jobs are running, interrupted, or completed, run the QUERY JOB command. For instructions, see QUERY JOB (Query a job). 3. To query retention rules, run the QUERY RETRULE command. For instructions, see QUERY RETRULE (Query a retention rule). 4. To query retention sets, run the QUERY RETSET command. For instructions, see QUERY RETSET (Query a retention set). 5. To query retention set contents, run the QUERY RETSETCONTENTS command. For instructions, see QUERY RETSETCONTENTS (Query the contents of a retention set).

Table 15. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information
<p>12 Monitor storage rules.</p>	<p>To obtain the overall status of storage rule operations, view the Storage Rules area on the Operations Center Overview page.</p>	<p>The status summary shows the most recent processing results for storage rules. The number of storage rules in each of the following states is shown:</p> <p> Normal The number of storage rules that ran without errors.</p> <p> Warning The number of storage rules that completed processing, but did not move or copy all eligible data. Either some files were skipped, the rule's time limit was reached, or the process was canceled.</p> <p> Failed The number of storage rules that did not complete processing. For example, the server might fail to process data because the target storage pool has insufficient space or because the server cannot access the storage pool.</p> <p> Other states The number of storage rules in other states. The server on which the storage rule is defined might be unavailable to provide the data, or might be running an earlier version of IBM Storage Protect that does not support status. Status might not be applicable because the storage rule was not activated or not run.</p> <p>Tips:</p> <ul style="list-style-type: none"> • An icon is displayed only if one or more storage rules are in the corresponding state. To view more detailed information about each storage rule, click Storage Rules to open the Storage Rules page. • To determine which storage rule jobs are running or completed, run the QUERY JOB command. For instructions, see QUERY JOB (Query a job).

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 settings icon  and click **Command Builder**.

Table 16. Periodic monitoring tasks		
Task	Basic procedures	Advanced procedures and troubleshooting
Monitor system performance.	<p>Determine the length of time that is required for client backup operations:</p> <ol style="list-style-type: none"> 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). 5. If the duration of client backup operations is increasing and the reasons are not clear, investigate the cause. <p>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:</p> <ol style="list-style-type: none"> 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. 	<p>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.</p> <p>Look for performance bottlenecks. For instructions, see Identifying performance bottlenecks.</p> <p>For information about identifying and resolving other performance issues, see Performance.</p>

Table 16. Periodic monitoring tasks (continued)



Task	Basic procedures	Advanced procedures and troubleshooting
Determine the disk savings that are provided by data deduplication.	<ol style="list-style-type: none"> 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. 	<p>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:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings 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). 3. View the statistical report by issuing the QUERY DEDUPSTATS command. Follow the instructions in QUERY DEDUPSTATS (Query data deduplication statistics).
Verify that current backup files for device configuration and volume history information are saved.	<p>Access your storage locations to ensure that the files are available. The preferred method is to save the backup files to two locations.</p> <p>To locate the volume history and device configuration files, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings icon  and click Command Builder. 2. To locate the volume history and device configuration files, issue the following commands: <pre>query option volhistory</pre> <pre>query option devconfig</pre> 3. In the output, review the Option Setting column to find the file locations. <p>If a disaster occurs, both the volume history file and the device configuration file are required to restore the server database.</p>	

Table 16. Periodic monitoring tasks (continued)


Task	Basic procedures	Advanced procedures and troubleshooting
Determine whether sufficient space is available for the instance directory file system.	<p>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:</p> <ul style="list-style-type: none"> AIX To view available space in the file system, on the operating system command line, issue the following command: <pre>df -g instance_directory</pre> where <i>instance_directory</i> specifies the instance directory. Linux To view available space in the file system, on the operating system command line, issue the following command: <pre>df -h instance_directory</pre> where <i>instance_directory</i> specifies the instance directory. Windows In the Windows Explorer program, right-click the file system and click Properties. View the capacity information. <p>The preferred location of the instance directory depends on the operating system where the server is installed:</p> <ul style="list-style-type: none"> AIX Linux /home/tsminst1/tsminst1 Windows C:\tsminst1 <p>Tip: If you completed a planning worksheet, the location of the instance directory is recorded in the worksheet.</p>	

Table 16. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
Identify unexpected client activity.	<p>To monitor client activity to determine whether data volumes exceed expected amounts, complete the following steps:</p> <ol style="list-style-type: none"> 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. 	<p>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.</p> <p>Periodically review the SQL activity summary table, which contains statistics about client sessions. To compare current activity with previous activity, use an SQL SELECT statement. If the level of activity is significantly different from previous activity, it might indicate a ransomware attack.</p> <p>Periodically review the activity log. Look for ANE messages that indicate how many files were backed up and inspected. Compare current data deduplication rates with previous rates. If an unusually high number of files were backed up, or the rate of data deduplication unexpectedly drops to 0, it might indicate a ransomware attack.</p>

Table 16. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor storage pool growth over time.	<ol style="list-style-type: none"> 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. 	<p>Tips:</p> <ul style="list-style-type: none"> • To specify the time period 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, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. 3. Specify the duration in the Delay period for container reuse field. • To determine data deduplication performance for directory-container and cloud-container storage pools, use the GENERATE DEDUPSTATS command. • To view data deduplication statistics for a storage pool, complete the following steps: <ol style="list-style-type: none"> 1. On the Storage Pools page of the Operations Center, select the storage pool. 2. Click Details > Properties. <p>Alternatively, use the QUERY EXTENTUPDATES command to display information about updates to data extents in directory-container or cloud-container storage pools. The command output can help you determine which data extents are no longer referenced and which ones are eligible to be deleted from the system. In the output, monitor the number of data extents that are eligible to be deleted from the system. This metric has a direct correlation to the amount of free space that will be available within the container storage pool.</p> • To display the amount of physical space that is occupied by a file space after the removal of the data deduplication savings, use the select * from occupancy command. The command output includes the LOGICAL_MB value. LOGICAL_MB is the amount of space that is used by the file space.

Table 16. Periodic monitoring tasks (continued)		
Task	Basic procedures	Advanced procedures and troubleshooting
Evaluate the timing of client schedules. Ensure that the start and end times of client schedules meet your business needs.	<p>On the Operations Center Overview page, click Clients > Schedules.</p> <p>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.</p>	<p>Tip: You can receive a warning message if a client operation runs longer than expected. Complete the following steps:</p> <ol style="list-style-type: none"> 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.
Evaluate the timing of maintenance tasks. Ensure that the start and end times of maintenance tasks meet your business needs.	<p>On the Operations Center Overview page, click Servers > Maintenance.</p> <p>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.</p>	<p>Tip: If a maintenance task is running too long, change the start time or the maximum run time. Complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Overview page, hover over the settings 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).

Related information

[QUERY ACTLOG \(Query the activity log\)](#)

[UPDATE STGPPOOL \(Update a storage pool\)](#)

[QUERY EXTENTUPDATES \(Query updated data extents\)](#)

Verifying license compliance

Verify that your IBM Storage Protect 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 that you use to verify that your solution complies with license terms varies depending on the provisions of your IBM Storage Protect 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.

Note: The front-end capacity measurements are ignored when the following set of IBM products send data to Amazon S3 compatible target via IBM Storage Protect Server:

- IBM Storage Protect Plus

- IBM Storage Protect for Cloud
- IBM Storage Defender
- IBM Storage Fusion

Back-end capacity licensing

The back-end model determines license requirements based on the terabytes of data that are stored in primary storage pools and repositories.

Tips:

- To ensure the accuracy of front-end and back-end capacity estimates, install the most recent version of the client software on each client node.
- The front-end and back-end capacity information in the Operations Center is for planning and estimation purposes.

PVU licensing

The PVU model is based on the use of PVUs by server devices.


Important: The PVU calculations that are provided by IBM Storage Protect are considered estimates and are not legally binding. The PVU licensing information that is reported by IBM Storage Protect is not considered an acceptable substitute for the IBM License Metric Tool. The IBM License Metric Tool is designed to reflect actual usage. For example, after you install the IBM Storage Protect backup-archive client, the tool counts the client only after first usage. For more information about the IBM License Metric Tool, see [IBM License Metric Tool](#).


If you have questions or concerns about licensing requirements, contact your IBM Storage Protect software provider.

Procedure

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

Tip: The Operations Center provides an email report that summarizes front-end and back-end capacity usage. Reports can be sent automatically to one or more recipients regularly. To configure and manage email reports, click **Reports** on the Operations Center menu bar.

Option	Description
Front-end model	<p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>The front-end capacity estimate is displayed on the Front-end Usage page.</p> <p>b. If a value is displayed in the Not Reporting column, click the number to identify clients that did not report capacity usage.</p> <p>c. To estimate capacity for clients that did not report capacity usage, go to the following download site, which provides measuring tools and instructions: https://public.dhe.ibm.com/storage/tivoli-storage-management/front_end_capacity_measurement_tools</p> <p>To measure front-end capacity by script, complete the instructions in the most recently available licensing guide.</p> <p>d. Add the Operations Center estimate and any estimates that you obtained by using a script.</p> <p>e. Verify that the estimated capacity complies with your licensing agreement.</p>
Back-end model	<p>Restriction: If the source and target replication servers do not use the same policy settings, you cannot use the Operations Center to monitor back-end</p>

Option	Description
	<p>capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <ol style="list-style-type: none"> On the Operations Center menu bar, hover over the settings icon  and click Licensing. Click the Back-end tab. Verify that the estimated amount of data complies with your licensing agreement.
PVU model	For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model .

Tracking system status by using email reports

Set up the Operations Center to generate email reports that summarize system status. You can configure a mail server connection, change report settings, and optionally create custom reports.

Before you begin

Before you set up 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 SMTP server must be configured as an open mail relay. You must also ensure that the IBM Storage Protect 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.
- To set up email reports, you must have system privilege for the server.
- 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 that is associated with it. To specify an email address for an administrator, use the **EMAILADDRESS** parameter of the **UPDATE ADMIN** command.

About this task

You can configure the Operations Center to send a general operations report, a license compliance report, and one or more custom reports. You create custom reports by selecting a template from a set of commonly used report templates or by entering SQL SELECT statements to query managed servers.

Procedure

To set up and manage email reports, complete the following steps:

1. On the Operations Center menu bar, click **Reports**.
2. If an email server connection is not yet configured, click **Configure Mail Server** and complete the fields.

After you configure the mail server, the general operations report and license compliance report are enabled.

3. To change report settings, select a report, click **Details**, and update the form.
4. Optional: To add a custom report, click **+ Report**, and complete the fields.

Tip: To immediately run and send a report, select the report and click **Send**.

Results

Enabled reports are sent according to the specified settings.

Related information

[UPDATE ADMIN \(Update an administrator\)](#)

Part 4. Managing operations for a multisite disk solution

Use this information to manage operations for a multisite disk solution with IBM Storage Protect that includes a server and uses data deduplication for multiple locations.

Managing the Operations Center

The Operations Center provides web and mobile access to status information about the IBM Storage Protect 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 IBM Storage Protect 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.

Before you begin

Communication between the spoke server and the hub server must be secured by using the Transport Layer Security (TLS) protocol. To secure communication, add the certificate of the spoke server to the truststore file of 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, and secure SSL/TLS communication is not required, 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 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 IBM Storage Protect 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
```

Tip: If a server definition is deleted immediately after the server is removed from the monitored group, status information for the server can remain in the Operations Center indefinitely.

To avoid this issue, wait until the status collection interval passes before you delete the server definition. The status collection interval is shown on the Settings page of the Operations Center.

Starting and stopping the web server

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

Before you begin

Note: The **systemctl** commands requires root privileges. The commands must be executed as a root user or with the **sudo** command.

Procedure

1. Stop 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 command:

```
./stopserver.sh
```

- **Linux** Issue the following command:

```
systemctl stop opscenter.service
```

- **Windows** From the **Services** window, stop the **IBM Storage Protect Operations Center** service.

2. 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 command:

```
./startserver.sh
```

- **Linux** Issue the following commands:

Start the server:

```
systemctl start opscenter.service
```

Restart the server:

```
systemctl restart opscenter.service
```

Determine whether the server is running:

```
systemctl status opscenter.service
```

- **Windows** From the **Services** window, start the **IBM Storage Protect 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 IBM Storage Protect 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 use the Operations Center to remove the hub server of IBM Storage Protect, and configure another 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.



Warning: If you change the hub server, or make a new or different server into the hub, any custom reports created on the old hub will not be carried over to the new hub. And you must recreate all of the Operations Center Settings, including Custom Reports, Shared Links, and Updates.

Related tasks

Restarting the initial configuration wizard

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

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 IBM Storage Protect 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 MONITOREDSEVERGROUP ""
```

```
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 MONITOREDSEVERGROUP ""
```

- 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

Restarting the initial configuration wizard

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

Starting and stopping the web server

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

Protecting applications, virtual machines, and systems

The 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 data protection solution with IBM Storage Protect, 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:

- IBM Storage Protect Snapshot
- IBM Storage Protect for Databases
- IBM Storage Protect for Enterprise Resource Planning
- IBM Storage Protect for Mail
- IBM Storage Protect 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 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 client software and planning the installation” on page 93](#).
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 94](#).
3. Specify when to back up and archive client data. Follow the instructions in [“Scheduling backup and archive operations” on page 97](#).

4. To allow the client to connect to the server, register the client. Follow the instructions in [“Registering clients”](#) on page 98.
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 99.

Selecting the client 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 backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Storage Protect for Databases, IBM Storage Protect for Enterprise Resource Planning, IBM Storage Protect for Mail, and IBM Storage Protect for Virtual Environments. If you install a product for which the client acceptor 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 99 before you can use the client.

Goal	Product and description	Installation instructions
Protect a file server or workstation	The 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.	<ul style="list-style-type: none"> • Client environment requirements • Install UNIX and Linux backup-archive clients • Installing the Windows client for the first time
Protect applications with snapshot backup and restore capabilities	IBM Storage Protect Snapshot 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.	<ul style="list-style-type: none"> • Installing and upgrading for UNIX and Linux • Installing and upgrading for VMware • Installing and upgrading for Windows
Protect an email application on an HCL Domino server	IBM Storage Protect for Mail: Data Protection for HCL Domino automates data protection so that backups are completed without shutting down HCL Domino servers.	<ul style="list-style-type: none"> • Installation of Data Protection for HCL Domino on a UNIX, AIX, or Linux system (V7.1.0) • Installation of Data Protection for HCL Domino on a Windows system (V7.1.0)
Protect an email application on a Microsoft Exchange server	IBM Storage Protect 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 Data Protection for Microsoft Exchange Server

Goal	Product and description	Installation instructions
Protect a Db2 database	The application programming interface (API) of the backup-archive client can be used to back up Db2 data to the IBM Storage Protect server.	Installing the backup-archive clients (UNIX, Linux, and Windows)
Protect an IBM Informix® database	The API of the backup-archive client can be used to back up Informix data to the IBM Storage Protect server.	Installing the backup-archive clients (UNIX, Linux, and Windows)
Protect a Microsoft SQL database	IBM Storage Protect for Databases: Data Protection for Microsoft SQL Server protects Microsoft SQL data.	Installing Data Protection for SQL Server on Windows Server Core
Protect an Oracle database	IBM Storage Protect for Databases: Data Protection for Oracle protects Oracle data.	Data Protection for Oracle installation
Protect an SAP environment	IBM Storage Protect 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.	<ul style="list-style-type: none"> • Installing Data Protection for SAP for Db2 • Installing Data Protection for SAP for Oracle
Protect a virtual machine	<p>IBM Storage Protect for Virtual Environments provides protection that is tailored for Microsoft Hyper-V and VMware virtual environments. You can use IBM Storage Protect for Virtual Environments to create incremental forever backups that are stored on a centralized server, create backup policies, and restore virtual machines or individual files.</p> <p>Alternatively, use the 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.</p>	<ul style="list-style-type: none"> • Installing and upgrading Data Protection for Microsoft Hyper-V • Installing and upgrading Data Protection for VMware • Installing the backup-archive clients (UNIX, Linux, and Windows)

Tip: To use the client for space management, you can install IBM Storage Protect for Space Management or IBM Storage Protect HSM for Windows.

Specifying rules for backing up and archiving client data

Before you add a client, ensure that appropriate rules are specified for backup and archive operations for 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 97](#).
- 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”](#) on page 95.
2. If you need to make minor changes to meet data retention requirements, follow the instructions in [“Editing policies”](#) on page 96.
3. Optional: If you need to create policy domains or make extensive changes to policies to meet data retention requirements, see [Customizing policies](#).

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.

Tip: To help ensure that you can recover data after a ransomware attack, apply the following guidelines:

- Ensure that the value in the Backups column is a minimum of 2. The preferred value is 3, 4, or more.
- Ensure that the value in the Keep Extra Backups column is a minimum of 14 days. The preferred value is 30 or more days.
- Ensure that the value in the Keep Archives column is a minimum of 30 days.

If IBM Storage Protect for Space Management software is installed on the client, ensure that data is backed up before you migrate it. On the **DEFINE MGMTCLASS** or **UPDATE MGMTCLASS** command, specify **MIGREQUIRESBKUP=YES**. Then, follow the guidelines in the tip.

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](#).

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. 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	<ol style="list-style-type: none"> a. In the Policy Sets table, click +Management Class. b. To specify the rules for backing up and archiving data, complete the fields in the Add Management Class window. c. To make the management class the default management class, select the Make default check box. d. Click Add.
Delete a management class	<p>In the Management Class column, click -.</p> <p>Tip: To delete the default management class, you must first assign a different management class as the default.</p>
Make a management class the default management class	<p>In the Default column for the management class, click the radio button.</p> <p>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.</p>

Option	Description
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:
- 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.
 - 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.
 - 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 98](#).
- 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. 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

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

Option	Description
Modify a schedule	a. In the Schedules view, select the schedule and click Details .

Option	Description
	b. On the Schedule Details page, view details by clicking the blue arrows at the beginning of the rows. c. 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 settings 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 more information about the commands, see [DEFINE SCHEDULE \(Define a client schedule\)](#) or [UPDATE SCHEDULE \(Update a client schedule\)](#).

Related information

[Tuning the schedule for daily operations](#)

Registering clients

Register a client to ensure that the client can connect to the server, and the server can protect client data.

Before you begin

Determine whether the client requires an administrative user ID with client owner authority over the client node. To determine which clients require an administrative user ID, see [technote 7048963](#).

Restriction: For some types of clients, the client node name and the administrative user ID must match. You cannot authenticate those clients by using the Lightweight Directory Access Protocol authentication method that was introduced in V7.1.7. For details about this authentication method, sometimes referred to as integrated mode, see [Authenticating users by using an Active Directory database](#).

Procedure

To register a client, complete one of the following actions.

- If the client requires an administrative user ID, register the client by using the **REGISTER NODE** command and specify the **USERID** parameter:

```
register node node_name password userid=node_name
```

where *node_name* specifies the node name and *password* specifies the node password. For details, see [Register a node](#).

- If the client does not require an administrative user ID, register the client by using the Operations Center Add Client wizard. Complete the following steps:
 - a. On the Operations Center menu bar, click **Clients**.
 - b. In the Clients table, click **+Client**.
 - c. Complete the steps in the **Add Client** wizard:
 - i) Specify that redundant data can be eliminated on both the client and server. In the Client-side data deduplication area, select the **Enable** check box.
 - ii) 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.

- iii) Follow the instructions in the wizard to specify the policy domain, schedule, and option set.
- iv) Set how risks are displayed for the client by specifying the at-risk setting.
- v) Click **Add Client**.

Related information

[Tcpserveraddress option](#)

[Tcpport option](#)

[Nodename option](#)

[Deduplication option](#)

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 100.

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

Software	Link to instructions
IBM Storage Protect backup-archive client	<ul style="list-style-type: none">– Install UNIX and Linux backup-archive clients– Installing the Windows client for the first time <p>Tip: You can also update existing clients by using the Operations Center. For instructions, see Scheduling client updates.</p>
IBM Storage Protect for Databases	<ul style="list-style-type: none">– Data Protection for Oracle installation– Installing Data Protection for SQL Server on Windows Server Core
IBM Storage Protect for Mail	<ul style="list-style-type: none">– Installation of Data Protection for HCL Domino on a UNIX, AIX, or Linux system (V7.1.0)– Installation of Data Protection for HCL Domino on a Windows system (V7.1.0)– Installing, upgrading, and migrating Data Protection for Microsoft Exchange Server
IBM Storage Protect Snapshot	<ul style="list-style-type: none">– Installing and upgrading for UNIX and Linux– Installing and upgrading for VMware– Installing and upgrading for Windows
IBM Storage Protect for Enterprise Resource Planning	<ul style="list-style-type: none">– Installing Data Protection for SAP for Db2– Installing Data Protection for SAP for Oracle

- 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 IBM Storage Protect backup-archive client.	<ul style="list-style-type: none"> – Install UNIX and Linux backup-archive clients – Installing the Windows client for the first time
If you plan to create incremental forever full backups of virtual machines, install and configure IBM Storage Protect for Virtual Environments and the backup-archive client on the same client node or on different client nodes.	<ul style="list-style-type: none"> – Data protection for VMware <p>Tip: You can obtain the software for IBM Storage Protect for Virtual Environments and the backup-archive client in the IBM Storage Protect for Virtual Environments installation package.</p>

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 98).

- For clients that are installed on an AIX, Linux, or Mac OS X 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 57.
4. Configure the client to run scheduled operations. Follow the instructions in [“Configuring the client to run scheduled operations”](#) on page 100.
5. Optional: Configure communications through a firewall. Follow the instructions in [“Configuring client/server communications through a firewall”](#) on page 102.
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

To change what is getting backed up from the client, follow the instructions in [“Modifying the scope of a client backup”](#) on page 106.

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 method is to install the backup-archive client on all client nodes so that you can configure and start the client acceptor on the client node. The client acceptor is designed to efficiently run scheduled operations. The client acceptor 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 client acceptor, you can reduce the number of background processes on the client and help to avoid memory retention problems.

The client acceptor runs schedules for the following products: the backup-archive client, IBM Storage Protect for Databases, IBM Storage Protect for Enterprise Resource Planning, IBM Storage Protect for Mail, and IBM Storage Protect for Virtual Environments. If you installed a product for which the client acceptor 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 client acceptor. Typically, third-party scheduling tools start 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 client acceptor, follow the instructions for the operating system that is installed on the client node:

AIX and Oracle Solaris

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

```
dsmc query sess
```

- Start the client acceptor by issuing the following command on the command line:

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

- To enable the client acceptor 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 # Client Acceptor Daemon
```

Linux

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

```
dsmc query sess
```

- Start the client acceptor by logging in with the root user ID and issuing the following command:

```
service dsmcad start
```

- g. To enable the client acceptor 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

- In the backup-archive client GUI, click **Edit > Client Preferences**.
- To ensure that the scheduler can start unattended, click **Authorization**, select **Password Generate**, and click **Apply**.
- To specify how services are managed, click **Web Client**, select **Schedule**, click **Apply**, and click **OK**.
- To ensure that the generated password is saved, restart the backup-archive client.
- Use the IBM Storage Protect Tools for Administrators application to start the client acceptor.

Windows

- In the backup-archive client GUI, click **Utilities > Setup Wizard > Help me configure the Client Scheduler**. Click **Next**.
- Read the information on the **Scheduler Wizard** page and click **Next**.
- On the **Scheduler Task** page, select **Install a new or additional scheduler** and click **Next**.
- On the **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**.
- Enter the name that you want to assign to this client acceptor. The default name is Client Acceptor. Click **Next**.
- Complete the configuration by stepping through the wizard.
- Update the client options file, `dsm.opt`, and set the **passwordaccess** option to generate.
- To store the client node password, issue the following command at the command prompt:

```
dsmc query sess
```

Enter the client node password when prompted.

- Start the client acceptor service from the **Services Control** page. For example, if you used the default name, start the Client Acceptor service. Do not start the scheduler service that you specified on the **Scheduler Name and Location** page. The scheduler service is started and stopped automatically by the client acceptor 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.

Before you begin

If you used the Add Client wizard to register a client, find the option values in the client options file that you obtained during that process. You can use the values to specify ports.

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 IBM Storage Protect 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 [KEEPALIVE](#).

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. The **tcpport** option in the client options 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.

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. 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.

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. 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 client problems](#).

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 57](#). For instructions about verifying the installation, see [“Verifying that the client management service is installed correctly” on page 58](#).

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:
 - a) On the Operations Center Overview page, click **Clients** and select the client.
 - b) Click **Details**.
 - c) On the client Summary page, click the **Diagnosis** tab.

- d) 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.

- e) 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

If you change the configuration of your solution, you must restart the client acceptor 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 client acceptor. The client acceptor 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 server, you must restart the client acceptor.

Procedure

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

AIX and Oracle Solaris

- To stop the client acceptor, complete the following steps:
 - a. Determine the process ID for the client acceptor by issuing the following command on the command line:

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

Review the output. In the following sample output, 6764 is the process ID for the client acceptor:

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

- b. Issue the following command on the command line:

```
kill -9 PID
```

where *PID* specifies the process ID for the client acceptor.

- To start the client acceptor, issue the following command on the command line:

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

Linux

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

```
# service dsmcad stop
```

- To stop and restart the client acceptor, issue the following command:

```
# service dsmcad restart
```

MAC OS X

Click **Applications > Utilities > Terminal**.

- To stop the client acceptor, issue the following command:

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

- To start the client acceptor, issue the following command:

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

Windows

- To stop the client acceptor service, complete the following steps:
 - a. Click **Start > Administrative Tools > Services**.
 - b. Double-click the client acceptor service.
 - c. Click **Stop** and **OK**.
- To restart the client acceptor service, complete the following steps:
 - a. Click **Start > Administrative Tools > Services**.
 - b. Double-click the client acceptor service.
 - c. Click **Start** and **OK**.

Related information

[Resolving client scheduling problems](#)

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 \(Unlock an administrator\)](#).
 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 \(Unlock a client node\)](#).
 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](#).

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](#).

- 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 option](#).
- 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 [Server-Client Compatibility and Upgrade Considerations](#). 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 [Supported Operating Systems](#).
3. If the solution includes storage agents or library clients, review the information about storage-agent and library-client compatibility with servers that are configured as library managers. See [Storage-agent and library-client compatibility with an server](#).

If you plan to upgrade a library manager and a library client, you must upgrade the library manager first.

Procedure

To upgrade the software, complete the instructions that are listed in the following table.

Software	Link to instructions
IBM Storage Protect backup-archive client	<ul style="list-style-type: none">• Scheduling client updates
IBM Storage Protect Snapshot	<ul style="list-style-type: none">• Installing and upgrading for UNIX and Linux• Installing and upgrading for VMware• Installing and upgrading for Windows
IBM Storage Protect for Databases	<ul style="list-style-type: none">• Upgrading Data Protection for SQL Server• Data Protection for Oracle installation
IBM Storage Protect for Enterprise Resource Planning	<ul style="list-style-type: none">• Upgrading Data Protection for SAP for Db2• Upgrading Data Protection for SAP for Oracle
IBM Storage Protect for Mail	<ul style="list-style-type: none">• Installation of Data Protection for HCL Domino on a Windows system (V7.1.0)• Installing, upgrading, and migrating Data Protection for Microsoft Exchange Server
IBM Storage Protect for Virtual Environments	<ul style="list-style-type: none">• Installing and upgrading Data Protection for VMware• Installing and upgrading Data Protection for Microsoft Hyper-V

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 the IBM Storage Protect 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:

- IBM Storage Protect Snapshot
- IBM Storage Protect for Databases
- IBM Storage Protect for Enterprise Resource Planning
- IBM Storage Protect for Mail
- IBM Storage Protect 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 file space.

Restriction: You cannot decommission an object client node.

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:
 - a) On the Operations Center **Overview** page, click **Clients** and select the client.
 - b) Click **More > Decommission**.
- To decommission a client node by using an administrative command, complete the following steps:
 - a) 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.

- b) 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
```

- c) 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 data center node is AUSTIN and the filespace ID is 7, issue the following command:

```
decommission vm austin 7 nametype=fsid
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example, if the virtual machine name is CODY 2 and the filespace name is \VMFULL - CODY 2, issue the following command:

```
decommission vm austin "\vmfull-cody 2"
```

- 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
```

If the virtual machine name includes one or more spaces, enclose the name in double quotation marks. For example, if the virtual machine name is CODY 2 and the filespace name is \VMFULL - CODY 2, issue the following command:

```
decommission vm austin "\vmfull-cody 2" wait=yes
```

What to do next

Watch for error messages, which might be displayed in the user interface or in the command output, immediately after you run the process.

You can verify that the client node is decommissioned:

1. On the Operations Center **Overview** page, click **Clients**.
2. In the Clients table, in the At risk column, review the state:
 - A DECOMMISSIONED state specifies that the node is decommissioned.
 - A null value specifies that the node is not decommissioned.
 - A PENDING state specifies that the node is being decommissioned, or the decommission process failed.

Tip: If you want to determine the status of a pending decommission process, issue the following command:

```
query process
```

3. Review the command output:

- If status is provided for the decommission process, the process is in progress. For example:

query process		
Process Number	Process Description	Process Status
-----	-----	-----
3	DECOMMISSION NODE	Number of backup objects deactivated for node NODE1: 8 objects deactivated.

- If no status is provided for the decommission process, and you did not receive an error message, the process is incomplete. A process can be incomplete if files that are associated with the node are not yet deactivated. After the files are deactivated, run the decommission process again.
- If no status is provided for the decommission process, and you receive an error message, the process failed. Run the decommission process again.

Related information

[DECOMMISSION NODE \(Decommission a client node\)](#)

[DECOMMISSION VM \(Decommission a virtual machine\)](#)

[QUERY NODE \(Query nodes\)](#)

[REMOVE REPLNODE \(Remove a client node from replication\)](#)

Deactivating data to free storage space

In some cases, you can deactivate data that is stored on the IBM Storage Protect 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 information

[DEACTIVATE DATA \(Deactivate data for a client node\)](#)

Managing data storage

Manage your data for efficiency and add supported devices and media to the server to store client data.

Related information

[Storage pool types](#)

Auditing a storage pool container

Audit a storage pool container to check for inconsistencies between database information and a container in a storage pool.

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

1. To audit a storage pool container, issue the **AUDIT CONTAINER** command.
For example, issue the following command to audit a container, 0000000000000076c.dcf:

```
audit container c:\tss-storage\07\0000000000000076c.dcf
```

2. Review the output from the ANR4891I message for 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 STGPPOOL** command.

Restriction: You can repair the contents of the storage pool only if you protected the storage pool by using the **PROTECT STGPPOOL** command.

Related information

[AUDIT CONTAINER](#) (Verify the consistency of database information for a directory-container storage pool)

[QUERY DAMAGED](#) (Query damaged data in a directory-container or cloud-container storage pool)

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 sizing the active log for your system, see [Planning the storage arrays](#).
- 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

- 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 [Resolving and preventing issues related to database growth and degraded performance in Tivoli Storage Manager V7.1.1.200 and later servers](#).
- To decrease the size of the database for V7.1 servers and later, issue the following IBM 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 REPLTLSPACE1 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
```

- To increase or decrease the size of the active log, complete the following steps:
 - 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.
 - 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 is based on the value of the **ACTIVELOGSIZE** option. Guidelines for space requirements are in the following table:

<i>Table 17. How to estimate volume and file space requirements</i>	
ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
16 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.
- 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 information

[ACTIVELOGSIZE server option](#)

[EXTEND DBSPACE \(Increase space for the database\)](#)

[SETOPT \(Set a server option for dynamic update\)](#)

Managing memory and processor usage

Ensure that you manage memory requirements and processor usage to ensure that the 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 required hardware and software. For more information, see [Supported Operating Systems](#).
- For more information about managing resources such as the database and recovery log, see [Planning the storage arrays](#).
- 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 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 65](#).
2. Optional: If the monitoring information shows that the server workload increased, review the planning information. Review whether the capacity of the system is adequate in the following cases:
 - The number of clients increases
 - The amount of data that is being backed up increases
 - The amount of time that is available for backups changes
3. Determine whether your solution is performing at the level you expect. Review the client schedules 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 **Backed up** and **Replicated** 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 tasks

[Defining schedules for server maintenance activities](#)

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

Related information

[Deduplicating data \(V7.1.1\)](#)

[Performance](#)

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 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. Allow enough space for future growth of workloads.

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 backups to the new server.

Procedure

To move a client node to another server, complete the following steps:

1. Replicate the client node data to a target replication server by using replication storage rules or export the client node directly to a new server by using the **EXPORT NODE** command.

Tip: Replicating data by using replication storage rules has some advantages over using the **EXPORT NODE** command. For example, replication storage rules help to support enhanced protection of data and provide the ability to configure replication operations with a high degree of flexibility and granularity. For more information, see [“Replicating client data by using replication storage rules” on page 117](#).

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 checkbox in the schedule row to which you want to assign the selected client node.
 - d. Click **Save**.
4. Run the replication storage rule or issue the **EXPORT NODE** command again to incrementally copy data from the original server to the new server. By incrementally copying data, you copy data that was backed up between the first export process and the time 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 information

[EXPORT NODE \(Export client node information\)](#)

[IMPORT NODE \(Import client node information\)](#)

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 the node level. You can also protect data at the storage-pool level.

Replication compatibility

Before you set up replication operations with IBM Storage Protect, you must ensure that the source and target replication servers are compatible for replication.

<i>Table 18. Replication compatibility of server versions</i>	
Source replication server version	Compatible versions for the target replication server
7.1	7.1 or later
7.1.1	7.1 or later
7.1.3	7.1.3 or later
7.1.4	7.1.3 or later
7.1.5	7.1.3 or later
7.1.6	7.1.3 or later
7.1.7	7.1.3 or later
7.1.8	7.1.3 or later
8.1	7.1.3 or later
8.1.1	7.1.3 or later
8.1.2	7.1.3 or later
8.1.3	7.1.3 or later
8.1.4	7.1.3 or later
8.1.5	7.1.3 or later
8.1.6	7.1.3 or later
8.1.7	7.1.3 or later
8.1.8	8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.9, 7.1.8, and 7.1.7
8.1.9	8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.9, 7.1.8, and 7.1.7
8.1.10	8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.9, 7.1.8, and 7.1.7
8.1.11	8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.9, 7.1.8, and 7.1.7
8.1.12	8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7

Table 18. Replication compatibility of server versions (continued)

Source replication server version	Compatible versions for the target replication server
8.1.13	8.1.13, 8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7
8.1.14	8.1.14, 8.1.13, 8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7
8.1.15	8.1.15, 8.1.14, 8.1.13, 8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7
8.1.16	8.1.16, 8.1.15, 8.1.14, 8.1.13, 8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7
8.1.17	8.1.17, 8.1.16, 8.1.15, 8.1.14, 8.1.13, 8.1.12, 8.1.11, 8.1.10, 8.1.9, 8.1.8, 8.1.7, 8.1.6, 8.1.1, 7.1.13, 7.1.12, 7.1.11, 7.1.10, 7.1.9, 7.1.8, and 7.1.7

Replicating client data by using replication storage rules

You can replicate data from a source replication server to two or three target replication servers. The process of creating copies of data and sending the copies to multiple servers is called *multi-target replication*. Replication of data helps to support disaster recovery and data availability if the source replication server becomes unavailable. The replication operation not only copies the data to the target replication server but also communicates the metadata, such as attributes of the source server nodes and file spaces.

About this task

Multi-target replication requires one source replication server and two or three target replication servers. After the servers are defined, to schedule multi-target replication, you must define multiple *replication storage rules*. Replication storage rules are designed for scheduling and automating the replication process. For example, to send data from a source replication server to three target replication servers, you would define three replication storage rules. By defining multiple replication storage rules, you can schedule, configure, and run the rules independently.

When you define a replication storage rule, you must specify the name of the target replication server. You can specify other parameters to define the starting time and maximum duration of the replication process. You can also define *subrules* for a replication storage rule. Subrules are designed to provide more granularity for defining the nodes, file spaces, and the data types to be replicated. When you define a subrule, you must specify the name of the replication storage rule, also referred to as the parent rule. When the replication storage rule runs, client node data is replicated. During subsequent processing of the replication storage rule, only new or changed client node data is replicated.

For users of the REPLICATE NODE command: When you transition from using the **REPLICATE NODE** command to using replication storage rules, only data that was previously not replicated or has changed is processed when you run replication storage rules.

Procedure

To define a storage rule for replication, follow the instructions in *Defining a replication storage rule* in IBM Documentation.

Enabling node replication

You can enable node replication to protect your data.

Before you begin

Ensure that the source and target servers are compatible for replication.

About this task

Replicate the client node to replicate all client data, including metadata. By default, node replication is disabled when you start the server for the first time.

Tips:

- To reduce replication processing time, protect the storage pool before you replicate client nodes. When node replication is started, the data extents that are already replicated through storage pool protection are skipped.
- Replication requires increased amounts of memory and 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:

- a) On the **Servers** page, click **Details**.
- b) On the **Details** page, click **Properties**.
- c) In the **Replication** section, select **Enabled** in the **Outbound replication** field.
- d) Click **Save**.

What to do next

Optionally, take one or both of the following actions:

- To verify that replication was successful, review the [“Daily monitoring checklist” on page 65](#).
- **Linux** If the IBM Storage Protect server replicates nodes to a remote server, determine whether Aspera® Fast Adaptive Secure Protocol (FASP®) technology can improve data throughput to the remote server. Follow the instructions in [Determining whether Aspera FASP technology can optimize data transfer in your system environment](#).

Related reference

[Replication compatibility](#)

Before you set up replication operations with IBM Storage Protect, you must ensure that the source and target replication servers are compatible for replication.

Protecting data in directory-container storage pools

Protect data in directory-container storage pools to reduce node replication time and to enable repair of data in directory-container storage pools.

Before you begin

Ensure that at least one directory-container storage pool exists on the target replication server. When you enable replication in the Operations Center, you can schedule storage pool protection. To configure replication and enable storage pool protection, complete the following steps:

1. On the Operations Center menu bar, hover over **Storage** and click **Replication**.
2. On the Replication page, click **Server Pair**.
3. Complete the steps in the Add Server Pair wizard.

About this task

Protecting a directory-container storage pool backs up data extents to another storage pool, and can improve performance for node replication. When node replication is started, the data extents that are already backed up through storage pool protection are skipped, which reduces the replication processing time. You can schedule the protection of storage pools several times a day to keep up with changes to data.

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.

Alternative protection strategy: As an alternative to using replication, you can protect data in directory-container storage pools by copying the data to container-copy storage pools. Data in container-copy storage pools is stored on tape volumes. Tape copies that are stored offsite provide additional disaster recovery protection in a replicated environment.

Procedure

1. Alternatively, to enable storage pool protection, you can use the **PROTECT STGPOOL** command from the source replication 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
```

As part of the operation of the **PROTECT STGPOOL** command, damaged extents in the target storage pool are repaired. To be repaired, extents must already be marked as damaged on the target replication server. For example, an **AUDIT CONTAINER** command might identify damage in the target storage pool before the **PROTECT STGPOOL** command is issued.

2. Optional: If damaged extents were repaired in the target storage pool and you protect multiple source storage pools in one target storage pool, complete the following steps to ensure a complete repair:
 - a) Issue the **PROTECT STGPOOL** command for all source storage pools to repair as much of the damage as possible.
 - b) Issue the **PROTECT STGPOOL** command again for all source storage pools. For this second operation, use the **FORCERECONCILE=YES** parameter.
This step ensures that any repairs from other source pools are properly recognized for all source storage pools.

Results

If a directory-container storage pool is protected, you can repair the storage pool if damage occurs, by using the **REPAIR STGPOOL** command.

Restriction: If you replicate client nodes but do not protect the directory-container storage pool, you cannot repair the storage pool.

What to do next

Optionally, take one or both of the following actions:

- To view replication workload status, follow the instructions in the [“Daily monitoring checklist”](#) on page 65.
- **Linux** If the IBM Storage Protect server replicates nodes to a remote server, determine whether Aspera Fast Adaptive Secure Protocol (FASP) technology can improve data throughput to the remote server. Follow the instructions in [Determining whether Aspera FASP technology can optimize data transfer in your system environment](#).

Related information

[Repairing and recovering data in directory-container storage pools](#)

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
- Changes to replication rules
- Requirement for a different server to be the target replication server
- Scheduled processes that negatively affect server performance

Procedure

Use the Operations Center to modify replication settings.

Task	Procedure
Change a replication rule.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.In the Replication section, choose the replication rule that you want to apply: Default archive rule, Default backup rule, or Default space-management rule.Click Save.
Specify the duration that replication records are retained.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.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.Click Save.
Specify a target replication server.	<ol style="list-style-type: none">On the Servers page, click Details.On the Details page, click Properties.In the Replication section, specify the target replication server.Click Save.
Cancel a replication process.	<ol style="list-style-type: none">On the Servers page, click Active tasks.Select the process or session that you want to cancel.Click Cancel.

Setting different retention policies for the source replication server and target replication server

You can set policies on the target replication server that manage the replicated client-node data differently than on the source replication server. For example, you can maintain a different number of versions of files on the source and the target replication servers.

Procedure

1. From the source replication server, validate the replication configuration and verify that the source replication server can communicate with the target replication server by issuing the **VALIDATE REPLICATION** command.
For example, validate the configuration by using the name of one client node that is being replicated:

```
validate replication node1 verifyconnection=yes
```

2. From the source replication server, issue the **VALIDATE REPLPOLICY** command to review the differences between the policies on the source and target replication servers.
For example, to display the differences between the policies on the source replication server and the target replication server, CVT_SRV2, issue the following command from the source replication server:

```
validate replpolicy cvt_srv2
```

3. Update the policies on the target replication server if necessary.

Tip: You can use the Operations Center to modify the policies on the target server. Follow the instructions in [“Editing policies” on page 96](#).

For example, to maintain inactive versions of files for a shorter time on the target replication server than on the source replication server, reduce the **Backups** setting in the management classes that apply to replicated client data.

4. Enable the target replication server to use its policies to manage the replicated client-node data by issuing the **SET DISSIMILARPOLICIES** command on the source replication server.
For example, to enable the policies on the target replication server, CVT_SRV2, issue the following command on the source replication server:

```
set dissimilarpolicies cvt_srv2 on
```

The next time that the replication process runs, the policies on the target replication server are used to manage the replicated client-node data.

Tip: If you configure replication by using the Operations Center and the policies on the source and target replication servers do not match, the policy that is specified for the source replication server is used. If you enabled the policies on the target replication server by using the **SET DISSIMILARPOLICIES** command, the policy that is specified for the target replication server is used. If the target replication server does not have the policy that is used by the node on the source replication server, the STANDARD policy is used.

Related information

[EXPORT POLICY \(Export policy information\)](#)

[SET DISSIMILARPOLICIES \(Enable the policies on the target replication server to manage replicated data\)](#)

[VALIDATE REPLICATION \(Validate replication for a client node\)](#)

[VALIDATE REPLPOLICY \(Verify the policies on the target replication server\)](#)

Replicating client node data after a database restore

When you restore the IBM Storage Protect database on a source replication server, replication is automatically disabled. Before re-enabling replication, you can take steps to preserve the client node data that is on the target replication server.

About this task

Disabling replication prevents the IBM Storage Protect server from deleting copies of data on the target replication server that are not referenced by the restored database. Before re-enabling replication, determine whether copies of data that are on the target replication server are needed. If they are, complete the steps that are described in the following example. In the example, the name of the source replication server is PRODSRV. DRSRV is the name of the target replication server. NODE1 is a client node with replicated data on PRODSRV and DRSRV.

Procedure

1. Remove NODE1 from replication on PRODSRV and DRSRV by issuing the **REMOVE REPLNODE** command:

```
remove replnode node1
```

2. Update NODE1 definitions PRODSRV and DRSRV. When replication occurs, DRSRV sends the data to PRODSRV that was lost because of the database restore.

- a. On DRSRV, issue the **UPDATE NODE** command and specify the replication mode **SYNCSEND**:

```
update node node1 replstate=enabled replmode=syncsend
```

- b. On PRODSRV, issue the **UPDATE NODE** command and specify the replication mode **SYNCRECEIVE**:

```
update node node1 replstate=enabled replmode=syncreceive
```

3. On DRSRV, set the replication rules to match the rules on PRODSRV. For example, if only archive data was being replicated from PRODSRV to DRSRV, set the rules on DRSRV to replicate only archive data from DRSRV to PRODSRV. Backup and space-managed data are not replicated to PRODSRV. To set rules, you can issue the following commands:

- UPDATE FILESPACE
- UPDATE NODE
- SET ARREPLRULEDEFAULT
- SET BKREPLRULEDEFAULT
- SET SPREPLRULE

4. On DRSRV, issue the **SET REPLSERVER** command to set PRODSRV as the target replication server:

```
set replserver prodsrv
```

5. On DRSRV, issue the **REPLICATE NODE** command to replicate data that is stored on NODE1:

```
replicate node node1
```

Replication processing changes the replication state of NODE1 to SEND on DRSRV and to RECEIVE on PRODSRV.

6. Remove NODE1 from replication on PRODSRV and DRSRV by issuing the **REMOVE REPLNODE** command:

```
remove replnode node1
```

7. Update NODE1 definitions:

- a. On DRSRV, issue the **UPDATE NODE** command and specify the replication mode SYNCRECEIVE:

```
update node node1 replstate=enabled replmode=syncreceive
```

- b. On PRODSRV, issue the **UPDATE NODE** command and specify the replication mode SYNCSEND:

```
update node node1 replstate=enabled replmode=syncsend
```

8. On PRODSRV, enable replication by issuing the **ENABLE REPLICATION** command:

```
enable replication
```

9. On PRODSRV, issue the **REPLICATE NODE** command to replicate data that is stored on NODE1:

```
replicate node node1
```

Replication processing changes the replication state of NODE1 to SEND on PRODSRV and to RECEIVE on DRSRV.

Results

The original replication configuration is restored. PRODSRV has all the data that was lost because of the database restore.

What to do next

Remember: In step 4, you set the PRODSRV as the target replication server for DRSRV. If in your original configuration you were replicating data from DRSRV to another server, you must reset the target replication server on DRSRV. For example, if you were replicating data from DRSRV to BKUPDRSRV, issue the following command on DRSRV:

```
set replserver bkupdrsrv
```

Securing the server

Secure the IBM Storage Protect 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 IBM Storage Protect 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 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 IBM Storage Protect 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 IBM Storage Protect server that the server, client, and storage agent use.

Restriction: Do not use the SSL or TLS protocols for communications with an IBM Db2 database instance that is used by any IBM Storage Protect servers.

Each 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 IBM Storage Protect 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 IBM Storage Protect server, client, and storage agent.

Authority levels

With each IBM Storage Protect server, different administrative authority levels are available that determine the 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 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.

Passwords

By default, the server automatically uses password authentication. With password authentication, all users must enter a password when they access the server.

Use Lightweight Directory Access Protocol (LDAP) to apply stricter requirements for passwords. For more information, see [Authenticating users by using an LDAP server](#).

Table 19. 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 the password is not changed within this period, the password must be changed the next time that the user accesses 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.
Default password length	The administrator can specify a minimum length for the password. Beginning with version 8.1.16, the default minimum length for server passwords changed from 8 to 15 characters and the minimum value that can be specified for the password length has changed from 1 to 8.

Session security

Session security is the level of security that is used for communication among IBM Storage Protect client nodes, administrative clients, and servers and is set by using the **SESSIONSECURITY** parameter.

The **SESSIONSECURITY** parameter can be set to one of the following values:

- The STRICT value enforces the highest level of security for communication between IBM Storage Protect servers, nodes, and administrators.
- The TRANSITIONAL value specifies that the existing communication protocol is used while you update your IBM Storage Protect software to V8.1.2 or later. This is the default. When **SESSIONSECURITY=TRANSITIONAL**, stricter security settings are automatically enforced as higher versions of the TLS protocol are used and as the software is updated to V8.1.2 or later. After a node, administrator, or server meets the requirements for the STRICT value, session security is automatically updated to the STRICT value, and the entity can no longer authenticate by using a previous version of the client or earlier TLS protocols.

Note: You are not required to update backup-archive clients to V8.1.2 or later before you upgrade servers. After you upgrade a server to V8.1.2 or later, nodes and administrators that are using earlier versions of the software will continue to communicate with the server by using the TRANSITIONAL value until the entity meets the requirements for the STRICT value. Similarly, you can upgrade backup-archive clients to V8.1.2 or later before you upgrade your IBM Storage Protect servers, but you are not required to upgrade servers first. Communication between servers and clients is not interrupted.

For more information about the **SESSIONSECURITY** parameter values, see the following commands.

Table 20. Commands used to set the SESSIONSECURITY parameter	
Entity	Command
Client nodes	<ul style="list-style-type: none"> • REGISTER NODE • UPDATE NODE
Administrators	<ul style="list-style-type: none"> • REGISTER ADMIN • UPDATE ADMIN
Servers	<ul style="list-style-type: none"> • DEFINE SERVER • UPDATE SERVER

Administrators that authenticate by using the **DSMADMC** command, **DSMC** command, or dsm program cannot authenticate by using an earlier version after authenticating by using V8.1.2 or later. To resolve authentication issues for administrators, see the following tips:

Tips:

- Ensure that all IBM Storage Protect software that the administrator account uses to log on is upgraded to V8.1.2 or later. If an administrator account logs on from multiple systems, ensure that the server's certificate is installed on each system.
- After an administrator successfully authenticates with the server by using V8.1.2 or later software or V7.1.8 or later software, the administrator can no longer authenticate with that server using client or server versions earlier than V8.1.2 or V7.1.8. An administrator command can be issued from any system.
- If necessary, create a separate administrator account to use only with clients and servers that are using V8.1.1 or earlier software.

Enforce the highest level of security for communication with the IBM Storage Protect server by ensuring that all nodes, administrators, and servers use STRICT session security. You can use the **SELECT** command to determine which servers, nodes, and administrators are using TRANSITIONAL session security and should be updated to use STRICT session security.

Related information

[Securing communications](#)

Managing administrators

An administrator who has system authority can complete any task with the IBM Storage Protect server, including assigning authority levels to other administrators. To complete some 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.	<p>To add an administrator, ADMIN1, with system authority and specify a password, complete the following steps:</p> <ol style="list-style-type: none">Register the administrator and specify Pa\$#\$tw0 as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$tw0</pre>Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>
Change administrative authority.	<p>Change the authority level for an administrator, ADMIN1.</p> <ul style="list-style-type: none">Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>Revoke system authority for the administrator by issuing the following command: <pre>revoke authority admin1 classes=system</pre>
Remove administrators.	<p>Remove an administrator, ADMIN1, from accessing the IBM Storage Protect server by issuing the following command: <pre>remove admin admin1</pre></p>
Temporarily prevent access to the server.	<p>Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.</p>

Changing password requirements

You can change the minimum password limit, password length, password expiration, and enable or disable authentication for IBM Storage Protect.

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 IBM Storage Protect servers.

Table 21. Authentication tasks for IBM Storage Protect servers	
Task	Procedure
Set a limit for invalid password attempts.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details, and then click the Properties tab. Set the number of invalid attempts in the Invalid sign-on attempt limit field. The default value at installation is 0.
Set a minimum length for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of characters in the Minimum password length field.
Set the expiration period for passwords.	<ol style="list-style-type: none"> On the Servers page in the Operations Center, select the server. Click Details and then click the Properties tab. Set the number of days in the Password common expiration field.
Set a default authentication method.	<p>Issue the SET DEFAULTAUTHENTICATION command. For example, to use the server as the default authentication method, issue the following command:</p> <pre>set defaultauthentication local</pre> <p>To update one client node to authenticate with the server, include AUTHENTICATION=LOCAL in the UPDATE NODE command:</p> <pre>update node authentication=local</pre>
Set the password complexity for the administrator accounts with SESSIONSECURITY=STRICT .	<p>To set the password complexity, issue the following commands and specify the integer values: SET MINPWCHARUPPER, SET MINPWCHARLOWER, SET MINPWCHARNUMERIC, and SET MINPWCHARSPECIAL.</p> <p>The sum of all the integer values that are specified by using these commands must be less than or equal to 58. Also, this sum of integer values is used as the minimum password length if it is greater than the value of integer that is specified by using the SET MINPWLENGTH command.</p>

Table 21. Authentication tasks for IBM Storage Protect servers (continued)

Task	Procedure
Set the password complexity for the administrator accounts with SESSIONSECURITY=TRANSITIONAL .	<p>To set the password complexity, issue the following commands and specify the integer values: SET MINPWCHARALPHABETIC, SET MINPWCHARNUMERIC, and SET MINPWCHARSPECIAL.</p> <p>The sum of all the integer values that are specified by using these commands must be less than or equal to 58. Also, this sum of integer values is used as the minimum password length if it is greater than the value of integer that is specified by using the SET MINPWLENGTH command.</p>
Set up multifactor authentication (MFA)	You can set up multifactor authentication (MFA) on an IBM Storage Protect server administrator account to provide an additional layer of protection. For more information, see <i>Configuring multifactor authentication</i> in IBM Documentation.

Related information

[Authenticating IBM Storage Protect users by using an LDAP server](#)

Securing IBM Storage Protect on the system

Protect the system where the IBM Storage Protect server runs to prevent unauthorized access.

Procedure

Ensure that unauthorized users cannot access the directories for the server database and the server instance. 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 IBM Storage Protect 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

- 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 126](#).
- To control the authority of an administrator to complete some tasks, use the following two server options:
 - 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.
 - 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.
- 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 the server by applying port restrictions.

About this task

You might have to restrict access to specific servers, based on your security requirements. The IBM Storage Protect server can be configured to listen on four TCP/IP ports: two that can be used for either regular TCP/IP protocols or Secure Sockets Layer (SSL)/Transport Layer Security (TLS) protocols and two that can be used only for the SSL/TLS protocol.

Procedure

You can set the server options to specify the port that you require, as listed in [Table 22 on page 129](#).

Table 22. Server options and port access	
Server option	Port access
TCPPORT	Specifies the port number on which the server TCP/IP communication driver is to wait for requests for client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default value is 1500.
TCPADMINPORT	<p>Specifies the port number on which the server TCP/IP communication driver is to wait for requests for sessions other than client sessions. This port listens for both TCP/IP and SSL-enabled sessions. The default is the value of TCPPORT.</p> <p>Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.</p>
SSLTCPPORT	Specifies the SSL TCP/IP port address for a server. This port listens for SSL-enabled sessions only. A default port value is not available.
SSLTCPADMINPORT	<p>Specifies the port address on which the server TCP/IP communication driver waits for requests for SSL-enabled sessions. A default port value is not available.</p> <p>Use this option to separate administrative client traffic from regular client traffic that uses the TCPPORT and SSLTCPPORT options.</p>

Restrictions:

The following restrictions apply when you specify the SSL-only server ports (**SSLTCPPORT** and **SSLTCPADMINPORT**):

- When you specify the server's SSL-only port for the **LLADDRESS** on the **DEFINE SERVER** or **UPDATE SERVER** command, you must also specify the **SSL=YES** parameter.
- When you specify the server's SSL-only port for the client's **TCPPORT** option, you must also specify **YES** for the SSL client option.

Related reference

[Planning firewall access](#)

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

Stopping and starting the server

Before you complete maintenance or reconfiguration tasks, stop the 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 IBM Storage Protect 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 or processes are in progress by completing the following steps:
 - a. On the **Overview** page of the Operations Center, view the **Activity** area for the total numbers of processes and sessions that are currently active. If numbers differ significantly from the usual numbers that are displayed during your daily storage-management routine, view other status indicators in the Operations Center to check whether there is a problem.
 - b. View the graph in the **Activity** area to compare the amount of network traffic over the following periods:
 - The current period, that is, the most recent 24-hour period
 - The previous period, that is, the 24 hours before the current periodIf the graph for the previous period represents the expected amount of traffic, significant differences on the graph for the current period might indicate a problem.
 - c. On the **Servers** page, select a server for which you want to view processes and sessions, and click **Details**. If the server is not registered as a hub or spoke server in the Operations Center, obtain information about processes by using administrative commands. Issue the **QUERY PROCESS** command to query processes and obtain information about sessions by issuing the **QUERY SESSION** command.
3. Wait until the client node sessions are completed or cancel them. To cancel processes and sessions, complete the following steps:

- On the **Servers** page, select a server for which you want to view processes and sessions, and click **Details**.
- Click the **Active Tasks** tab, and select one or more processes, sessions, or a combination of both that you want to cancel.
- Click **Cancel**.
- If the server is not registered as a hub or spoke server in the Operations Center, cancel sessions by using administrative commands. Issue the **CANCEL SESSION** command to cancel a session and cancel processes by using the **CANCEL PROCESS** command.

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.

4. 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
```

Tip: To view a video about starting the server in maintenance mode, see [Starting a server in maintenance mode](#).

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. For instructions, see [Starting the server instance](#). Follow the instructions for your operating system.

Operations that were disabled during maintenance mode are reenabled.

Planning to upgrade the server

When a fix pack or interim fix becomes available, you can upgrade the IBM Storage Protect server to take advantage of product improvements. 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 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 [Fix Central](#).
2. Review product improvements, which are described in readme files.
Tip: When you obtain the installation package file from the [support site](#), you can also access the readme file.
3. Ensure that the version that you upgrade your server to is compatible with other components, such as storage agents and library clients. See [Storage-agent and library-client compatibility with an server](#).
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 [Server-Client Compatibility and Upgrade Considerations](#).
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 in [Installing an server fix pack](#).

Preparing for an outage or system update

Prepare IBM Storage Protect to maintain your system in a consistent state during a planned 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. In the Operations Center, 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
```

Tip: You can issue the halt command from the Operations Center by hovering over the **Settings** icon and clicking **Command Builder**. Then, select the server, type `halt`, and press **Enter**.

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 a recovery server. Use replication and storage pool protection to protect data. You must also determine how often directory-container storage pools are protected.

Completing recovery drills

Schedule disaster recovery drills to prepare for audits that certify the recoverability of the IBM Storage Protect 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 replication server at recovery site and recovery time is fast. When there is an outage, the source replication server can automatically fail over to a target replication server for data recovery. If a disaster occurs and the source replication 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

[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 the node level. You can also protect data at the storage-pool level.

Related information

[Replicating client node data after a database restore \(V7.1.1\)](#)

Recovering from data loss or system outages

You can use IBM Storage Protect to recover data that was lost when a disaster or system outage occurred. You can recover directory-container storage pools, client data, and databases.

Before you begin

Schedule client and server workloads to achieve the best performance for your storage environment. Issue the **PROTECT STGPOOL** and **REPLICATE NODE** commands as part of the schedule. 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.

Procedure

Use the following recovery methods based on the component that you must recover.

Component to recover	Procedure	More information
Directory-container storage pool	<p>To recover directory-container storage pools, complete the following steps:</p> <ol style="list-style-type: none">Scan for damaged data extents in the directory-container storage pool by using the AUDIT CONTAINER command and specifying the ACTION=SCANALL parameter.Repair damaged data extents in the directory-container storage pool by using the REPAIR STGPOOL command. <p>Restriction: You can repair a storage pool only if the storage pool is protected.</p> <ol style="list-style-type: none">Remove damaged data extents by using the AUDIT CONTAINER command and specifying the ACTION=REMOVEDAMAGED parameter.	<p>“Repairing storage pools” on page 137</p>

Component to recover	Procedure	More information
Client data	<p>Prerequisites:</p> <ul style="list-style-type: none"> • 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. <p>Manually configure the client to automatically fail over to the target replication server for data recovery.</p> <p>If you enabled the client for automated client failover, you can recover the data by using automatic failover function. You can verify that the <code>usereplicationfailover</code> option is either not in the client options file or is set to yes. Recover data from the target replication server when the source replication server is unavailable due to an outage by using automatic failover.</p> <p>Tip:</p> <ul style="list-style-type: none"> • Use the SET FAILOVERHLADDRESS command to specify the IP address for the replication server during failover, if the address is different from the IP address that is specified for the replication process. 	<p>SET FAILOVERHLADDRESS (Set a failover high level address)</p>
Database	<p>Prerequisites:</p> <ul style="list-style-type: none"> • 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. <p>Restore the IBM Storage Protect database to the most current state or to a specific point in time by using the DSMSERV RESTORE DB server utility.</p>	<p>DSMSERV RESTORE DB (Restore the database)</p>

Related information

[AUDIT CONTAINER \(Verify the consistency of database information for a directory-container storage pool\)](#)
[DSMSERV RESTORE DB \(Restore the database\)](#)

Restoring the database

You might have to restore the IBM Storage Protect database after a disaster. You can restore the database to the most current state or to a specified point in time. You must have full, incremental, or snapshot database backup volumes to restore the database to the current state or to a specified point in time. If you do not have the required backup volumes, you can still restore the database by re-creating the required backup files.

Before you begin

Restrictions:

- To restore the database to its latest version, you must locate the archive log directory. If you cannot locate the directory, you can restore the database only to a point in time.
- If the database and recovery log directories are lost, you can re-create the directories. However, you must first re-create the server instance. For detailed instructions about re-creating files and directories that are required to restore a server instance, see *Restoring a server when files are missing* in IBM Documentation.
- If the release level of the database backup is different from the release level of the server that is being restored, you cannot restore the server database. For example, if you are using a version 8.1 server and you try to restore a version 7.1 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. Depending on the version of the database that you want to restore, choose one of the following methods:

- 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 2022, use the following command:

```
dsmserv restore db todote=04/19/2022
```

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 to display a list of containers:

- a. From an administrative client, issue the following command:

```
select container_name from containers
```

- b. From the file system, issue the following command for the storage pool directory on the source replication server:

Tip: The storage pool directory is displayed in the command output:

```
Linux | AIX [root@source]$ ls -lR
```

```
Windows c:\source_stgpoldir>dir /s
```

- 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.

Tip: After a database restore operation, if containers exist on the file system that are not referenced in the server database, the **QUERY STGPOOL** command inaccurately displays the storage pool usage. When you restore a database to a point in time, containers might remain on the file system, but are not referenced in the server database. To help ensure accurate statistics about storage pool usage, you must manually delete any containers that are available on the file system, but are not referenced in the server database.

Related information

[Synchronization of source and target replication servers after role reversal](#)

[AUDIT CONTAINER \(Verify the consistency of database information for a directory-container storage pool\)](#)

[DSMSERV RESTORE DB \(Restore the database\)](#)

Repairing storage pools

If a disaster or system outage occurs, 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.

If you used the **REPLICATE NODE** command to replicate data, you can issue the **REPAIR STGPOOL** command for a specified storage pool only if the storage pool is protected by using the **PROTECT STGPOOL** command.

If you used replication storage rules to replicate data from a source replication server to one or more target replication servers, you do not have to issue the **PROTECT STGPOOL** command. This type of replication combines the functionality of the **REPLICATE NODE** and **PROTECT STGPOOL** commands into one replication operation.

Restriction: For data extents to be repairable, the data must have been replicated to a container storage pool on the target replication server. If data was replicated to a non-container storage pool or data was tiered out of the container storage pool on the target replication server, those data extents will not be recoverable.

Procedure

1. Repair the directory-container storage pool by taking one of the following actions:

- If you replicated data by using the **REPLICATE NODE** command, issue the **REPAIR STGPOOL** command.

For example, to repair a storage pool that is named STGPOOL1, issue the following command:

```
repair stgpool stgpool1
```

- If you used replication storage rules to replicate data, issue the **REPAIR STGPOOL** command and specify the name of the target replication server for the **SERVER** parameter.

For example, to repair a storage pool that is named STGPOOL2 from a replication server that is named SERVER2, issue the following command:

```
repair stgpool stgpool2 server=server2
```

2. If you replicated data by using the **REPLICATE NODE** command, complete this step:
If the damaged storage pool is specified as a target storage pool on the **PROTECT STGPOOL** command for one or more source storage pools, issue the **PROTECT STGPOOL** command for all source storage pools.
3. To ensure that all damaged data is identified and repaired from other source storage pools, take one of the following actions:
 - If you replicated data by using the **REPLICATE NODE** command, issue the **PROTECT STGPOOL** command again from all source storage pools and specify the **FORCERECONCILE=YES** parameter.
 - If you replicated data by using one or more replication storage rules, run the replication storage rule again to repair damaged extents in the storage pool.
4. To remove objects that refer to damaged data, issue the **AUDIT CONTAINER** command and specify the **ACTION=REMOVEDAMAGED** parameter.
5. If the damaged storage pool is a target storage pool for node replication from one or more source servers, take one of the following actions:
 - If you replicated data by using the **REPLICATE NODE** command, issue the **REPLICATE NODE** command again from all source servers.
 - If you replicated data by using one or more replication storage rules, run the replication storage rule from all source servers.
6. If you replicated data by using the **REPLICATE NODE** command, complete this step:
When the damage is repaired, issue the **PROTECT STGPOOL** command to ensure that the storage pool is protected to another directory-container storage pool.

What to do next

Ensure that no damaged data extents are displayed in the output by using the **QUERY DAMAGED** command.

Related information

[Repairing and recovering data in directory-container storage pools](#)

[AUDIT CONTAINER \(Verify the consistency of database information for a directory-container storage pool\)](#)

[QUERY DAMAGED \(Query damaged data in a directory-container or cloud-container storage pool\)](#)

[REPAIR STGPOOL \(Repair a directory-container storage pool\)](#)

Synchronization of source and target replication servers after role reversal

If a source replication server fails or becomes temporarily unavailable, you might have to reverse the server roles and temporarily configure the target replication server as the source replication server. Later, when the original source replication server is back online, you can restore the original roles. Then, you can replicate all data from the target replication server to the source replication server and synchronize the data on both servers. If you configure multi-target replication by using replication storage rules, you have an additional option to help ensure data availability and recovery if the source replication server fails. When you replicate data to two target replication servers, you can convert one of the target replication servers to a source replication server. The data from the newly established source replication server can then be replicated to another target replication server by using replication storage rules to help provide uninterrupted data protection.

Example for reversing server roles in a single-target replication environment

The following values are used in this example:

- NODE1 is the client node.
 - PRODSRV is the source replication server to which data from the client node, NODE1, is backed up.
 - DRSRV is the target replication server.
1. The PRODSRV server becomes unavailable. The system administrator, Elizabeth, updates the node definition on the DRSRV server to revoke PRODSRV server's role as the target of backup data for NODE1. She follows the instructions in "REMOVE REPLNODE (Remove a client node from replication)" in IBM Documentation and issues the following command on the DRSRV server:

```
remove replnode node1 server=prodsrv
```

By setting NODE1 as a non-replicating node, DRSRV can now be used as a source replication server for backing up the client data.

2. Elizabeth configures NODE1 to back up data to the DRSRV server. To achieve this, she must update the appropriate client options file.

Elizabeth reconfigures the appropriate client options file by following the instructions in "Creating and modifying the client system-options file" in IBM Documentation and "Creating and modifying the client options file" in IBM Documentation.

All backup operations for NODE1 are now directed to the DRSRV server.

After some time, the PRODSRV server comes back online. Elizabeth intends to restore the original storage setup.

3. She updates the node definition for NODE1 on the DRSRV server to be a non-replicating node. She issues the following command on the PRODSRV server:

```
remove replnode node1 server=drsrv
```

4. To start the process of synchronizing the servers, Elizabeth follows the instructions in "UPDATE NODE (Update node attributes)" in IBM Documentation and issues the following commands.

She issues the following command on the DRSRV server:

```
update node node1 replstate=enabled replmode=syncsend
```

She issues the following command on the PRODSRV server:

```
update node node1 replstate=enabled replmode=syncreceive
```

After these commands are run, the inventories of both IBM Storage Protect servers will be synchronized during the next replication operation.

5. To replicate all data from the DRSRV server to the PRODSRV server, Elizabeth defines a replication storage rule. She follows the instructions in "DEFINE STGRULE (Define a storage rule for replicating data)" in IBM Documentation and issues the following command on the DRSRV server:

```
define stgrule repl_role_reversal prodsrv actiontype=replicate
```

By issuing the command, she establishes the PRODSRV server as the default target replication server for the DRSRV server.

6. To synchronize the servers, Elizabeth follows the instructions in "START STGRULE (Start a replication rule)" in IBM Documentation and issues the **START STGRULE** command to replicate the data.

She issues the following command on the DRSRV server:

```
start stgrule repl_role_reversal
```

After the replication operation, the node definition on the DRSRV server is set to **REPLMODE=SEND**. The node definition on the PRODSRV server is set to **REPLMODE=RECEIVE**.

Elizabeth waits for the replication operation to be completed successfully before proceeding to the next step. Waiting for completion of the replication operation is mandatory.

7. Elizabeth updates the node definitions on the PRODSRV and DRSRV servers to set them as non-replicating nodes.

She issues the following command on the DRSRV server:

```
remove replnode node1 server=prodsrv
```

She issues the following command on PRODSRV server:

```
remove replnode node1 server=drsrv
```

8. Elizabeth configures the NODE1 node to set PRODSRV as the source replication server and DRSRV as the target replication server. She synchronizes the nodes on both servers by issuing the following commands.

She issues the following command on the PRODSRV server:

```
update node node1 replstate=enabled replmode=syncsend
```

She issues the following command on the DRSRV server:

```
update node node1 replstate=enabled replmode=syncreceive
```

9. Elizabeth configures the NODE1 node to back up data to the PRODSRV server by updating the appropriate client options file.

She follows the instructions in "Creating and modifying the client system-options file" in IBM Documentation and "Creating and modifying the client options file" in IBM Documentation.

All backup operations for the NODE1 node are now directed to the PRODSRV server.

10. Elizabeth defines a replication storage rule to configure replication of data from the NODE1 node on the PRODSRV server to the DRSRV server.

She issues the following command on the PRODSRV server:

```
define stgrule repl_drsrv drsrv actiontype=replicate
```

11. She runs the replication storage rule, REPL_DRSRV. After the replication operation is completed successfully, the node definition on the PRODSRV server is set to **REPLMODE=SEND** and the node definition on the DRSRV server is set to **REPLMODE=RECEIVE**.

At this point, the two servers are in the original configuration and the PRODSRV server has the data that was stored to the DRSRV server while the PRODSRV server was unavailable.

Appendix A. Accessibility features for the IBM Storage Protect product family

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Storage Protect family of products includes the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Storage Protect family of products uses the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the product.

The product documentation in IBM Documentation is enabled for accessibility.

Keyboard navigation

This product uses standard navigation keys.

Interface information

User interfaces do not have content that flashes 2 - 55 times per second.

Web user interfaces rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

Web user interfaces include WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Storage Protect 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 accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility](http://www.ibm.com/able) (www.ibm.com/able).

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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*

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 jurisdictions 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 wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
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.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, OpenShift®, Ansible®, and Ceph® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

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 Storage Protect family of products.

See the [IBM Storage Protect glossary](#).

Index

A

- About this publication [vii](#)
- access
 - limit [129](#)
 - server options [129](#)
- accessibility features [141](#)
- active log capacity [111](#)
- archive log capacity [111](#)
- archive operations
 - scheduling [97](#)
 - specifying rules [94](#)
- Aspera FASP [118](#)
- Aspera Fast Adaptive Secure Protocol, *See* Aspera FASP
- AUDIT CONTAINER [111](#)
- audit storage pool [111](#)
- authority level [126](#)

B

- back-end capacity licensing [82](#)
- backup operations
 - modifying scope [106](#)
 - scheduling [97](#)
 - specifying rules [94](#)

C

- client acceptor
 - configuring [100](#)
 - restarting [104](#)
 - stopping [104](#)
- client management service
 - configure Operations Center to use [59](#)
 - installing [57](#)
 - verify installation [58](#)
- client nodes
 - decommissioning [108](#)
 - removing from production [108](#)
- client/server communications
 - configuring [102](#)
- clients
 - adding [92](#)
 - assign to schedules [56](#)
 - configuring [56, 99](#)
 - configuring to run scheduled operations [100](#)
 - connecting to server [98](#)
 - decommission [115](#)
 - define schedules [56](#)
 - installing [56, 99](#)
 - managing operations [103](#)
 - move [115](#)
 - protecting [92](#)
 - register [56](#)
 - registering [98](#)
 - selecting software [93](#)
 - upgrading [107](#)

- commands
 - HALT [130](#)
 - REPAIR STGPOOL [137](#)
- configuration
 - changing [104](#)
 - clients [99](#)
- configuring
 - clients [56](#)
 - spoke server [87](#)

D

- daily checklist of monitoring tasks [65](#)
- data
 - deactivating [110](#)
- data deduplication
 - configure [52](#)
- data loss [134](#)
- data recovery
 - strategy [133](#)
- data retention rules
 - define [53](#)
- database capacity [111](#)
- deactivation process
 - backup data [110](#)
- decommission process
 - client node [108](#)
- disability [141](#)
- disaster recovery [133](#)
- disaster recovery manager [133](#)
- DRM [133](#)

E

- email reports
 - configuring [84](#)
- error logs
 - evaluating [103](#)

F

- file systems
 - [preparing, AIX server systems [41](#)
 - planning for [10](#)
 - preparing, Linux server systems [42](#)
 - preparing, Windows server systems [43](#)
- firewall [24, 25](#)
- firewalls
 - configuring communications through [102](#)
- front-end capacity licensing [82](#)

G

- graphical wizard
 - prerequisite RPM files [45](#)

H

- halting
 - server [130](#)
- hardware requirements [6](#)
- hub server
 - changing [90](#)
 - restore to preconfigured state [90](#)
 - secure SSL communications [60](#)

I

- IBM Documentation [vii](#)
- IBM License Metric Tool [82](#)
- IBM Storage Protect directories
 - planning for [10](#)
- implementation
 - test operations [62](#)
- initial configuration wizard
 - configure [89](#)
- install server
 - AIX systems [44](#)
 - Linux systems [44](#)
 - Windows systems [45](#)
- installation
 - clients [99](#)
- installing
 - clients [56](#)
- installing the operating system
 - AIX server systems [28](#)
 - Linux server systems [30](#)
 - Windows server systems [35](#)
- inventory capacity [111](#)
- issues
 - diagnosing [65](#)

K

- keyboard [141](#)

L

- LDAP
 - password requirements [126](#)
- license compliance
 - verifying [82](#)

M

- maintenance
 - define schedule [53](#)
- maintenance mode
 - start server [130](#)
- maintenance tasks
 - scheduling [114](#)
 - start the server in maintenance mode [131](#)
- managing
 - access levels [128](#)
 - administrators [126](#)
 - authority [126](#)
- managing security [123](#)
- memory requirements
 - managing [113](#)

- monitoring
 - daily checklist [65](#)
 - goals [65](#)
 - periodic checklist [77](#)
 - tasks
 - daily checklist [65](#)
 - periodic checklist [77](#)
- multi-site replication [117](#)
- multi-target replication [117](#)
- multipath I/O
 - configure for AIX systems [35](#)
 - configure for Linux systems [36](#)
 - configure for Windows systems [37](#)
- multisite disk solution
 - planning for [1](#)

N

- node replication
 - enable [62](#)

O

- operating system
 - install on AIX server systems [28](#)
 - install on Linux server systems [30](#)
 - install on Windows server systems [35](#)
 - security [128](#)
- Operations Center
 - configure [49](#)
 - restore to preconfigured state [90](#)
 - secure communications [50](#)
 - spoke server [87](#)
 - web server [88](#)
- options
 - set for server [48](#)
- outage
 - prepare [132](#)

P

- password requirements
 - LDAP [126](#)
- passwords
 - changing [126](#)
 - resetting [105](#)
- periodic checklist of monitoring tasks [77](#)
- planning solutions
 - multisite disk [1](#)
- planning worksheet [10](#)
- policies
 - editing [96](#)
 - specifying [94](#)
 - viewing [95](#)
- policy domains
 - specifying [94](#)
- privilege class
 - system privilege [126](#)
- processor usage [113](#)
- processor value unit (PVU) licensing [82](#)
- product license
 - register [52](#)
- publications [vii](#)

R

- reconfiguration tasks
 - start the server in maintenance mode [131](#)
- recovery
 - disaster recovery [133](#)
 - strategy [133](#)
- recovery drill [133](#)
- recovery method
 - data loss [134](#)
 - system outage [134](#)
- registration
 - clients [98](#)
- repair storage pool
 - damaged [137](#)
- replication
 - enabling [118](#)
 - managing [116](#)
 - modifying [120](#)
 - multisite disk solution
 - compatibility [116](#)
 - target server policies [121](#)
- reports
 - email
 - configuring [84](#)
- restricting
 - user access [128](#)
- RPM files
 - install for graphical wizard [45](#)
- rules
 - editing [96](#)
 - specifying
 - backup and archive operations [94](#)
 - viewing [95](#)

S

- scheduled activities
 - tuning [114](#)
- schedules
 - backup and archive operations [97](#)
- second server
 - add as spoke [62](#)
 - configure [60](#)
- secure communications
 - configure with SSL and TLS [49](#)
- security [123](#)
- server
 - configure [46](#)
 - configure second server [60](#)
 - create user ID for [38](#)
 - define maintenance schedule [53](#)
 - determine size of [4](#)
 - enabling replication [118](#)
 - enabling replication target policies [121](#)
 - managing replication [116](#)
 - modifying replication [120](#)
 - node replication [118](#)
 - plan upgrade [132](#)
 - set options [48](#)
 - start in maintenance mode [130](#)
 - stop [130](#)
- server installation
 - AIX systems [44, 45](#)

- server installation (*continued*)
 - Linux systems [44, 45](#)
- servers
 - start in maintenance mode [131](#)
- shutting down
 - server [130](#)
- software
 - selecting [93](#)
- Software requirements
 - Linux [8](#)
- solution
 - expanding [92](#)
- spoke server
 - add [62](#)
 - adding [87](#)
 - remove [88](#)
- spoke servers
 - restore to preconfigured state [90](#)
- SSL [49](#)
- starting server
 - maintenance mode [130](#)
- status reports
 - obtaining [84](#)
- stopping
 - server [130](#)
- storage configuration
 - planning for [10](#)
- storage hardware
 - configure [28](#)
- storage pool
 - protection [118](#)
 - repair [118, 137](#)
- storage pools
 - auditing containers [111](#)
- storage space
 - releasing [110](#)
- system requirements
 - hardware [6](#)
- system status
 - tracking [84](#)
- system update
 - prepare [132](#)

T

- TLS [49](#)
- troubleshooting
 - administrator IDs [105](#)
 - errors in client operations [103](#)
 - locked client nodes [105](#)
 - password issues [105](#)

U

- upgrade
 - server [132](#)
- user ID
 - create for server [38](#)

W

- web server
 - start [88](#)

web server (*continued*)
stop [88](#)



Product Number: 5725-W98
5725-W99
5725-X15