

IBM Storage Protect
8.1.24

Tape Solution Guide



Note:

Before you use this information and the product it supports, read the information in [“Notices” on page 217](#).

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
Tape planning requirements.....	2
System requirements for a tape-based solution.....	3
Hardware requirements.....	3
Software requirements.....	6
Planning worksheets.....	8
Planning for disk storage.....	13
Planning the storage arrays.....	13
Planning for tape storage.....	15
Supported tape devices and libraries.....	15
Supported tape device configurations.....	16
Data movement between storage devices.....	16
Library sharing.....	17
LAN-free data movement.....	17
Mixed device types in libraries.....	18
Definitions for Tape Storage Devices.....	20
Planning the storage pool hierarchy.....	21
Offsite data storage.....	24
Planning for security.....	24
Planning for administrator roles.....	24
Planning for secure communications.....	25
Planning for storage of encrypted data.....	26
Planning firewall access.....	27
Part 2. Implementing.....	29
Setting up the system.....	30
Configuring the storage hardware.....	30
Installing the server operating system.....	31
Installing on AIX systems.....	31
Installing on Linux systems.....	32
Installing on Windows systems.....	37
Configuring multipath I/O.....	38
AIX systems.....	38
Linux systems.....	39
Windows systems.....	40
Creating the user ID and directories for the server instance.....	41
Preparing file systems for the server.....	43
AIX systems.....	43
Linux systems.....	45
Windows systems.....	46
Installing the server and Operations Center.....	46
Installing on AIX and Linux systems.....	46
Installing prerequisite RPM files for the graphical wizard.....	47
Installing on Windows systems.....	48

Configuring the server and the Operations Center.....	49
Configuring the server instance.....	49
Installing the backup-archive client.....	50
Setting options for the server.....	50
Security concepts.....	51
Configuring secure communications with Transport Layer Security.....	54
Configuring the Operations Center.....	54
Securing communications between the Operations Center and the hub server.....	55
Registering the product license.....	57
Defining data retention rules for your business.....	57
Defining schedules for server maintenance activities.....	58
Moving backup media.....	62
Moving retention set data to and from tape storage.....	67
Defining client schedules.....	75
Attaching tape devices for the server.....	75
Attaching an automated library device to your system.....	76
Setting the library mode.....	76
Selecting a tape device driver.....	76
IBM tape device drivers.....	77
IBM Storage Protect tape device drivers.....	77
Special file names for tape devices.....	78
Installing and configuring tape device drivers.....	79
Installing and configuring IBM device drivers for IBM tape devices.....	79
AIX systems.....	83
Linux systems.....	85
Windows systems.....	88
Configuring libraries for use by a server.....	90
Defining tape devices.....	91
Defining libraries and drives.....	91
Defining tape device classes.....	93
Configuring library sharing.....	100
Example: Library sharing for AIX and Linux servers.....	101
Example: Library sharing for Windows servers.....	102
Setting up a storage pool hierarchy.....	105
Protecting applications and systems.....	107
Adding clients.....	107
Selecting the client software and planning the installation.....	108
Specifying rules for backing up and archiving client data.....	109
Scheduling backup and archive operations.....	112
Registering clients.....	113
Installing and configuring clients.....	114
Configuring LAN-free data movement.....	118
Validating your LAN-free configuration.....	119
Encryption methods.....	120
Configuring tape drive encryption.....	122
Controlling tape storage operations.....	123
How IBM Storage Protect fills volumes.....	123
Specifying the estimated capacity of tape volumes.....	124
Specifying recording formats for tape media.....	124
Associating library objects with device classes.....	125
Controlling media-mount operations for tape devices.....	125
Controlling the number of simultaneously mounted volumes.....	125
Controlling the amount of time that a volume remains mounted.....	126
Controlling the amount of time that the server waits for a drive.....	126
Preempting operations.....	127
Mount point preemption.....	127
Volume access preemption.....	128
Impacts of device changes on the SAN.....	128

Displaying device information.....	129
Write-once, read-many tape media.....	129
WORM-capable drives.....	130
Check-in of WORM media.....	130
Restrictions on WORM media.....	130
Mount failures with WORM media.....	131
Relabeling WORM media.....	131
Removing private WORM volumes from a library.....	131
Creation of DLT WORM volumes.....	131
Support for short and normal 3592 WORM tapes.....	131
Querying a device class for the WORM-parameter setting.....	131
Troubleshooting problems with devices.....	131
Completing the implementation.....	133
Part 3. Monitoring.....	135
Daily checklist.....	135
Periodic checklist.....	145
Monitoring tape alert messages for hardware errors.....	151
Preventing errors caused by media incompatibility.....	151
Operations with cleaner cartridges.....	152
Verifying license compliance.....	152
Tracking system status by using email reports.....	154
Part 4. Managing.....	155
Managing the Operations Center.....	155
Managing client operations.....	155
Evaluating errors in client error logs.....	155
Stopping and restarting the client acceptor.....	156
Resetting passwords.....	157
Managing client upgrades.....	158
Decommissioning a client node.....	159
Deactivating data to free storage space.....	161
Managing data storage.....	161
Managing inventory capacity.....	162
Tuning scheduled activities.....	163
Optimizing operations by enabling collocation of client files.....	164
Effects of collocation on operations.....	166
Selecting volumes with collocation enabled.....	167
Selecting volumes with collocation disabled.....	169
Collocation settings.....	170
Collocation of copy storage pools.....	170
Collocation of retention storage pools.....	171
Planning for and enabling collocation.....	171
Managing tape devices.....	173
Preparing removable media.....	173
Labeling tape volumes.....	174
Checking storage volumes into a library.....	176
Managing volume inventory.....	181
Controlling access to volumes.....	181
Reusing tapes.....	181
Maintaining a supply of scratch volumes.....	183
Maintaining a supply of volumes in a library that contains WORM media.....	183
Manage the volume inventory in automated libraries.....	184
Partially written volumes.....	188
Shared library operations.....	188
Server requests for volumes.....	189
Managing tape drives.....	191

Updating drives.....	191
Taking tape drives offline.....	192
Data validation during read/write operations to tape.....	192
Supported drives.....	193
Enabling and disabling logical block protection.....	194
Read/write operations to volumes.....	195
Storage pool management in a tape library.....	195
Cleaning tape drives.....	196
Methods for cleaning tape drives.....	196
Configuring the server for drive cleaning in an automated library.....	197
Resolving errors that are related to drive cleaning.....	199
Tape drive replacement.....	199
Deleting tape drives.....	199
Replacing drives with others of the same type.....	200
Migrating data to upgraded drives.....	201
Securing the server.....	201
Managing administrators.....	201
Changing password requirements.....	202
Securing the server on the system.....	204
Restricting user access to the server.....	204
Stopping and starting the server.....	204
Stopping the server.....	205
Starting the server for maintenance or reconfiguration tasks.....	206
Planning to upgrade the server.....	206
Preparing for an outage.....	207
Preparing for and recovering from a disaster by using DRM.....	208
Disaster recovery plan file	208
Recovering the server and client data.....	210
Recovery drills.....	211
Restoring the database.....	212
Appendix A. Accessibility.....	215
Notices.....	217
Glossary.....	221
Index.....	223

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. Planning for a tape-based data protection solution

Plan for a data protection solution that includes disk-to-disk-to-tape and disk-to-tape backup operations to optimize storage.

Planning roadmap

Plan for the tape solution by reviewing the architecture layout in [Figure 1 on page 1](#) and then completing the roadmap tasks that follow the diagram.




Figure 1. Tape solution

In this data protection configuration, the server uses both disk and tape storage hardware. Storage pool staging is used, in which client data is initially stored in disk storage pools and then later migrated to tape storage pools. For disaster recovery, tape volumes can be stored offsite. Offsite options include physically moving a second copy offsite by a courier or electronically vaulting copies offsite to a remote library.

Tips:

- In the described solution, data is *migrated* from disk storage pools to tape storage pools. However, instead of migrating the data, you can use the tiering-to-tape feature that was introduced in IBM Storage Protect Version 8.1.8. With this feature, you can automatically tier data from directory-container storage pools on disk to tape storage. You can specify that all data is tiered based on a specified age threshold, or that only inactive data is tiered based on an age threshold. For more information about tiering data to tape storage, see [Tiering data to cloud, tape, or file storage](#).
- The described solution does not include node replication. If you want to use node replication to back up a storage pool from disk to disk, ensure that the replication operation is completed before data is migrated from disk to tape. You can also use node replication to back up a storage pool on a local tape device to a copy storage pool on a local tape device.

To plan for a tape-based solution, complete the following tasks:

1. [Meet system requirements for hardware and software.](#)
2. [Record values for your system configuration in the planning worksheets.](#)
3. [Plan for disk storage.](#)
4. [Plan for tape storage.](#)
5. [Plan for security.](#)

Tape planning requirements

Before you implement a tape solution, review the general guidelines about system requirements. Determine whether to back up data to disk or tape, or a combination of both.

Network bandwidth

The network must have sufficient bandwidth for the expected data transfers between the client and the server, and for the cross-site restore operations that are required for disaster recovery. Use a storage area network (SAN) for data transfers among the server, disk devices, and tape devices. For more information, see [“Hardware requirements” on page 3](#).

Data migration

Migrate all data from disk to tape daily. Specify a FILE device class for disk-based storage pools. Schedule migration to control when processing occurs. To prevent automatic migration based on the migration threshold, specify a value of 100 for the **HIGHMIG** parameter and 0 for the **LOWMIG** parameter when you issue the **DEFINE STGPOOL** command. You must keep at least 20% of the tape drives available for restore operations. To use up to 80% of available tape drives and improve throughput performance, specify the **MIGPROCESS** parameter.

Consider the following information based on the type of data that is migrated:

- Use tape to back up data from clients that have large objects, such as databases.
Tip: Check with your tape-drive manufacturer for guidance about the size of the database that is suitable to write to tape.
- Use disk to back up data from clients that have smaller objects.
- To back up data directly to tape, use LAN-free data movement. For more information, see [“Configuring LAN-free data movement” on page 118](#).
- Do not back up virtual machines to tape. Use a separate disk-based storage pool that does not migrate to a tape-based storage pool. For more information about virtual machine support, see [and IBM Tivoli Storage Manager guest support for Virtual Machines and Virtualization](#).

Storage pool capacity

Maintain enough storage pool capacity to allow for 2 days of client backups and a buffer of 20%. You might have to schedule full backups over a few days to ensure that you have enough storage pool space.

Tape drives

Review the manufacturer specifications and estimate the capacity of a tape drive. Determine the amount of space that is required for backup and migration operations. Reserve 20% of tape drives for restore operations.

Related information

[MIGRATE STGPOOL \(Migrate storage pool to next storage pool\)](#)

System requirements for a tape-based solution

Hardware and software requirements are provided for a tape-based storage solution that has a data ingestion rate of 14 TB per hour.

Review the information to determine the hardware and software requirements for your storage environment. You might have to make adjustments based on your system size.

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 more information about planning disk devices, see [Planning for disk storage](#).

For more information about planning tape devices, see [Planning for tape storage](#).

The following table includes minimum hardware requirements for the server and storage. If you are using local partitions (LPARs) or work partitions (WPARs), adjust the network requirements to take account of the partition sizes. The figures in the table are based on a data ingestion rate of 14 TB per hour.

Hardware component	System requirements
Server processor	<div><div>AIX</div>8 processor cores, 3.42 GHz or faster. For example, use a POWER8® processor-based server.</div> <div><div>Linux</div><div>Windows</div>16 processor cores, 2.0 GHz or faster. For example, use an Intel Xeon processor.</div>
Server memory	64 GB RAM.
Network	The following sizing manages approximately 14 TB of data per hour: <ul style="list-style-type: none">• 10 Gb Ethernet (a minimum of four ports)• 8 Gb Fibre Channel adapter (a minimum of four ports) The number of ports depends on the percentage of daily data ingestion to disk storage pools versus tape storage. Use separate Fibre Channel adapters for tape and disk data.

Hardware component	System requirements
Storage	<p>Disk</p> <p>Based on the amount of data that you are writing to disk, specify the number of disks that you require.</p> <p>Ensure that the sequential input/output (I/O) throughput of the storage area network (SAN) matches the I/O throughput for the network in the previous row.</p> <p>For example, if you must back up 10 TB of data in a four-hour window, the throughput is approximately 700 MB per second. In this case, the server requires a front-end network (client-to-server path) that supports a minimum throughput of 700 MB per second. The back-end SAN (the server-to-storage device path) also must support a minimum throughput of 700 MB per second.</p> <p>To calculate the required disk speed, use the following formulas:</p> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> $\frac{(\text{Total amount of daily data ingestion} - \text{amount of daily data ingestion directly to tape}) \div (\text{Number of hours for daily client backup operations})}{\text{Megabytes of data ingestion to disk per hour}}$ </div> <div style="background-color: #f0f0f0; padding: 10px; margin: 10px 0;"> $\frac{(\text{Megabytes of data ingestion to disk per hour}) \div (3600 \text{ seconds per hour})}{\text{Megabytes of data ingestion per second that must be supported by the disk technology}}$ </div> <p>Tape</p> <p>Select the tape technology that best fits your business requirements. For example, use IBM Linear Tape-Open (LTO) or IBM TS1150 tape drives. Ensure that you have sufficient mount points for client backup operations and for migration. For more information about planning tape storage, see Planning for tape storage. For a list of supported tape devices, see IBM Support Portal for IBM Storage Protect.</p> <p>Tip: To optimize data movement, use LAN-free data movement.</p>
SAN I/O adapters	<p>Segregate disk and tape I/O. For more information about selecting an adapter, see the documentation for Brocade hardware products and for IBM Storwize® storage solutions.</p> <p>Disk</p> <p>Use at least two adapters.</p> <p>Tape</p> <p>Use at least two adapters.</p>

Estimating 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 IBM Storage Protect 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 IBM Storage Protect server is a workload that requires extra space for database operations on both the hub server and any spoke servers. The amount of space on the hub server for the archive log is larger if the hub server is monitoring one or more spoke servers. Review the following guidelines to estimate how much space your IBM Storage Protect server requires.

Database space for the Operations Center

The Operations Center uses approximately 4.4 GB of database space for every 1000 clients that are monitored on that server. This calculation applies to both hub servers and spoke servers within a configuration.

For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1000 clients. This configuration has a total of 5000 clients across the four servers. Each of the spoke servers requires 4.4 GB of database space. If the spoke servers are at IBM Storage Protect Version 8.1.2 or later, the hub server requires 8.8 GB of database space for monitoring only its 2000 clients:

$$(4.4 \text{ GB} \times 2) = 8.8 \text{ GB}$$

Database space for managed data

Managed data is the amount of data that is protected, including the amount of data for all retained versions.

- For client types that perform incremental-forever backups, the following formula can be used to estimate the total managed data:

$$\text{Front-end} + (\text{front-end} \times \text{change rate} \times (\text{retention} - 1))$$

For example, if you back up 100 TB of front-end data, use a 30-day retention period, and have a 5% change rate, calculate your total managed data by using the following figures:

$$100 \text{ TB} + (100 \text{ TB} \times 0.05 \times (30-1)) = 245 \text{ TB total managed data}$$

- For client types that perform full backups every day, the following formula can be used to estimate the total managed data:

$$\text{Front-end} \times \text{retention} \times (1 + \text{change rate})$$

For example, if you back up 10 TB of front-end data, use a 30-day retention period, and have a 3% change rate, calculate your total managed data by using the following figures:

$$10 \text{ TB} \times 30 \times (1 + .03) = 309 \text{ TB total managed data}$$

Unstructured data, average object size: 4 MB

Structured data, average object size: 128 MB

Unstructured data, number of objects =

$$(245 \text{ TB} \times 1024 \times 1024) / 4 \text{ MB} = 64225280$$

Structured data, number of objects =

$$(309 \text{ TB} \times 1024 \times 1024) / 128 \text{ MB} = 2531328$$

Total number of objects: 66756608

Managed data cost (1 KB per object) =

$$(66756608 \text{ KB}) / (1024 \times 1024) = 63.66 \text{ GB}$$

Plan for 20% of additional space so that database systems are not at 100% capacity:

$$\text{Database total physical storage requirements} = (\text{managed data space} + \text{Operations Center space}) \times (1.20)$$

For this example, you would calculate the space by using the following figures:

$$(66.33 \text{ GB} + 8.4 \text{ GB}) \times 1.20 = 76.41 \text{ GB}$$

Archive log space

The Operations Center uses approximately 18 GB of archive log space every 24 hours, per server, for every 1000 clients monitored on that server. Additionally, for every 1000 clients that are monitored on spoke servers, additional archive log space is used on the hub server. For spoke servers at V8.1.2 or later, this added amount is 1.2 GB of archive log space on the hub server per 1000 clients monitored every 24 hours.

For example, consider a hub server with 2000 clients that also manages three spoke servers, each with 1000 clients. This configuration has a total of 5000 clients across the four servers. You can calculate the archive log space for the hub server by using the following formula:

$$((18 \text{ GB} \times 2) + (1.2 \text{ GB} \times 3)) = 39.6 \text{ GB of archive log space}$$

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 requirements with a collection interval of once every 3 minutes for a configuration in which V8.1.2 or later spoke servers are monitored:

- Hub server: In the range 39.6 GB - 66 GB
- Each spoke server: In the range 18 GB - 30 GB

Allocate archive log space so that you can support the Operations Center without affecting server operations.

Software requirements

Documentation for the IBM Storage Protect tape-based solution includes installation and configuration tasks for IBM AIX®, Linux®, and Microsoft Windows operating systems. You must meet the minimum software requirements that are listed.

AIX systems

Type of software	Minimum software requirements
Operating system	IBM AIX 7.2 <ul style="list-style-type: none">• AIX 7.2 TL4 with SP3, TL5 with SP2, or later AIX 7.2 levels IBM AIX 7.3 <ul style="list-style-type: none">• AIX 7.3. TL0 with SP1 or later AIX 7.3 levels For more information about operating system requirements, see the IBM Storage Protect installation information.
Gunzip utility	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.

Type of software	Minimum software requirements
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	Korn Shell (ksh)

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>

Type of software	Minimum software requirements
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	<p>Microsoft Windows Server 2019</p> <p>Microsoft Windows Server 2016</p>
File system type	NTFS
Other software	<p>Windows 2019 or Windows 2016 with .NET Framework 3.5 is installed and enabled.</p> <p>The following User Account Control policies must be disabled:</p> <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

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.

Table 1. Worksheet for preconfiguration of a server system				
Item	Default value	Your value	Minimum directory size	More information
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	<div>Linux AIX</div> /home/tsminst1/tsminst1 <div>Windows</div> C:\tsminst1		<div>AIX</div> 50 GB. <div>Linux Windows</div> 25 GB.	If you change the value for the server instance directory from the default, also modify the Db2 instance owner value in Table 2 on page 11 .
Directory for server installation	<ul style="list-style-type: none"><div>Linux AIX</div> /<div>Windows</div> C:		<div>AIX</div> Available space that is required for the directory: 5 GB. <div>Linux Windows</div> Minimum space that is required for the directory: 30 GB	
Directory for server installation	/usr		<div>AIX</div> Available space that is required for the directory: 5 GB.	
Directory for server installation	<div>AIX</div> /var		<div>AIX</div> Available space that is required for the directory: 5 GB.	
Directory for server installation	<div>AIX</div> /tmp		<div>AIX</div> Available space that is required for the directory: 5 GB.	

Table 1. Worksheet for preconfiguration of a server system (continued)

Item	Default value	Your value	Minimum directory size	More information
Directory for server installation	AIX /opt		AIX Available space that is required for the directory: 10 GB.	
Directory for the active log	Linux AIX / tsminst1/TSMalog Windows C:\tsminst1\TSMalog		128 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	Linux AIX / tsminst1/TSMarchlog Windows C:\tsminst1\TSMarchlog		3 TB.	
Directories for the database	Linux AIX /tsminst1/TSMdbspace00 /tsminst1/TSMdbspace01 /tsminst1/TSMdbspace02 /tsminst1/TSMdbspace03 Windows C:\tsminst1\TSMdbspace00 C:\tsminst1\TSMdbspace01 C:\tsminst1\TSMdbspace02 C:\tsminst1\TSMdbspace03		For instructions about calculating space requirements, see “Hardware requirements” on page 3.	Create four file systems for the database.
Directories for storage	Linux AIX /tsminst1/TSMfile00 /tsminst1/TSMfile01 /tsminst1/TSMfile02 /tsminst1/TSMfile03 ... Windows C:\tsminst1\TSMfile00 C:\tsminst1\TSMfile01 C:\tsminst1\TSMfile02 C:\tsminst1\TSMfile03 ...		Determine the minimum total capacity for all directories by using the following calculation: Daily percentage of ingested data that is written to disk + 20% = Minimum total capacity	The preferred method is to define at least one directory for each tape device.

Table 2. Worksheet for IBM Storage Protect configuration

Item	Default value	Your value	More information
Db2 instance owner	tsminst1		If you changed the value for the server instance directory in Table 1 on page 9 from the default, also modify the value for the Db2 instance owner.
Db2 instance owner password	<div>Linux</div> <div>AIX</div> <div>Windows</div> passwd 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	<div>Linux</div> <div>AIX</div> tsmsrvrs		
Server name	The default value for the server name is the system host name.		
Server password	passwd		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	passwd		Select a different value for the administrator password than the default. Ensure that you record this value in a secure location.

Table 2. Worksheet for IBM Storage Protect configuration (continued)

Item	Default value	Your value	More information
Schedule start time	23:00		<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> <p>In this guide, the suggested time to start client backup operations is 23:00.</p>

Table 3. Worksheet for tape configuration

Item	Default value	Your value	More information
Robotic device files	<p>IBM devices with an IBM tape device driver:</p> <ul style="list-style-type: none"> • AIX /dev/smcX • Linux /dev/IBMchangerX • Windows ChangerX <p>Non-IBM devices with an IBM Storage Protect device driver:</p> <ul style="list-style-type: none"> • AIX /dev/lbX • Linux /dev/tsmcsi/lbX • Windows lbA.B.C.D 		<p>To manually define the library device files, use the following commands:</p> <ul style="list-style-type: none"> • DEFINE LIBRARY • DEFINE DRIVE • DEFINE PATH <p>For SCSI, you can use the PERFORM LIBACTION command to define all drives and their paths for a single library in one step. To use this command to define all drives and paths, the SANDISCOVERY option must be supported and enabled.</p>

Table 3. Worksheet for tape configuration (continued)			
Item	Default value	Your value	More information
Tape drives	<p>IBM devices with an IBM tape device driver:</p> <ul style="list-style-type: none"> • AIX /dev/rmtX • Linux /dev/IBMtapeX • Windows TapeX <p>Non-IBM devices with an IBM Storage Protect device driver:</p> <ul style="list-style-type: none"> • AIX /dev/mtX • Linux /dev/tsm SCSI/mtX • Windows mtA.B.C.D 		

Planning for disk 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 about selecting the appropriate storage hardware and setup for your solution, review the following guidelines.

Database, active log, and archive log

- Use a solid-state disk (SSD) or a fast, 15,000 rpm disk for the IBM Storage Protect database and active log.
- When you create arrays for the database, use RAID level 5.
- Use separate disks for archive log and database backup storage.

Storage pool

Use RAID level 6 for storage pool arrays to add protection against double drive failures when you use large disk types.

Planning the storage arrays

Prepare for disk storage configuration by planning for RAID arrays and volumes, according to 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 disk 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 4 on page 14 and Table 5 on page 14 describe the layout requirements for storage groups and volume configuration.

<i>Table 4. Components of storage group configuration</i>	
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: Quantity x Usable size
Suggested storage group and array names	Preferred name to use for MDisks and MDisk groups.
Usage	Server component that uses part of the physical disk.

<i>Table 5. Components of volume configuration</i>	
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 tape storage

Determine which tape devices to use and how to configure them. To optimize system performance, plan to use fast, high-capacity tape devices. Provision enough tape drives to meet your business requirements.

Supported tape devices and libraries

The server can use a wide range of tape devices and libraries. Select tape devices and libraries that meet your business requirements.

For a list of supported devices and valid device class formats, see the website for your operating system:

- [AIX](#) | [Windows](#) [Supported devices for AIX and Windows](#)
- [Linux](#) [Supported devices for Linux](#)

For more information about storage devices and storage objects, see [Types of storage devices](#).

Each device that is defined to IBM Storage Protect is associated with one *device class*. The device class specifies the device type and media management information, such as recording format, estimated capacity, and labeling prefixes.

A *device type* identifies a device as a member of a group of devices that share similar media characteristics. For example, the LTO device type applies to all generations of LTO tape drives.

A device class for a tape drive must also specify a library. A *physical library* is a collection of one or more drives that share similar media-mounting requirements. That is, the drive can be mounted by an operator or by an automated mounting mechanism.

A *library object definition* specifies the library type and other characteristics that are associated with that library type.

The following table lists the preferred library types for an IBM Storage Protect Version 8.1.6 tape solution.

Table 6. Library types for an IBM Storage Protect 8.1.6 tape solution		
Library type	Description	More information
SCSI	<p>A SCSI library is controlled through a SCSI interface, attached either directly to the server's host by using SCSI cabling or by a storage area network. A robot or other mechanism automatically handles tape volume mounts and dismounts.</p> <p>If you create different drive types for a SCSI library, you create multiple logical libraries that cannot be split between different types of drives. A SCSI library can contain drives of mixed technologies, including LTO Ultrium and digital linear tape (DLT) drives. For example:</p> <ul style="list-style-type: none">• The Oracle StorageTek L700 library• The IBM 3592 tape device	<p>“Configuring libraries for use by a server” on page 90</p> <p>Restrictions apply when you mix different generations of media and drives. For more information, see:</p> <ul style="list-style-type: none">• “Mixing generations of 3592 drives and media in a single library” on page 97• “Mixing LTO drives and media in a library” on page 94

Table 6. Library types for an IBM Storage Protect 8.1.6 tape solution (continued)

Library type	Description	More information
Shared	<p>Shared libraries are logical libraries that are represented by SCSI. The library is controlled by the IBM Storage Protect server that is configured as a library manager.</p> <p>IBM Storage Protect servers that use the SHARED library type are library clients to the library manager server. Shared libraries reference a library manager.</p>	

Supported tape device configurations

Review the information about local area networks (LAN) and storage area networks (SAN). To optimize data movement, plan to configure LAN-free data movement. In addition, consider whether to use library sharing.

Select the device configuration that meets your business requirements.

LAN-based and LAN-free data movement

You can move data between clients and storage devices that are attached to a local area network (LAN), or to storage devices that are attached to a storage area network (SAN), known as LAN-free data movement.

In a conventional LAN configuration, one or more tape libraries are associated with a single IBM Storage Protect server. LAN-free data movement makes LAN bandwidth available for other uses and decreases the load on the IBM Storage Protect server.

In a LAN configuration, client data, email, terminal connection, application program, and device control information must be handled by the same network. Device control information and client backup and restore data flow across the LAN.

A SAN is a dedicated storage network that can improve system performance.

By using IBM Storage Protect in a SAN, you benefit from the following functions:

- Sharing storage devices among multiple IBM Storage Protect servers.
- **Restriction:** A storage device with the GENERICTAPE device type cannot be shared among servers.
- Moving IBM Storage Protect client data directly to storage devices (LAN-free data movement) by configuring a storage agent on the client system.

In a SAN, you can share tape drives and libraries that are supported by the IBM Storage Protect server, including most SCSI tape devices.

When IBM Storage Protect servers share a SCSI tape, one server, the *library manager*, owns and controls the device. The storage agents, along with other IBM Storage Protect servers that share this library are *library clients*. A library client requests shared library resources, such as drives or media, from the library manager, but uses the resources independently. The library manager coordinates the access to these resources. IBM Storage Protect servers that are defined as library clients use server-to-server communications to contact the library manager and request device service. Data moves over the SAN between each server and the storage device.

Requirement: If you define a library manager server that is shared with the IBM Storage Protect server, the **SANDISCOVERY** option must be set to ON. By default, this option is set to OFF.

IBM Storage Protect servers use the following features when sharing an automated library:

Partitioning of the volume inventory

The inventory of media volumes in the shared library is partitioned among servers. Either one server owns a particular volume, or the volume is in the global scratch pool. No server owns the scratch pool.

Serialized drive access

Only one server accesses each tape drive at a time. Drive access is serialized. IBM Storage Protect controls drive access so that servers do not dismount other servers' volumes or write to drives where other servers mount their volumes.

Serialized mount access

The library autochanger completes a single mount or dismount operation at a time. The library manager completes all mount operations to provide this serialization.

Library sharing

You can optimize the efficiency of your tape solution by configuring library sharing. Library sharing allows multiple IBM Storage Protect servers to use the same tape library and drives on a storage area network (SAN) and to improve backup and recovery performance and tape hardware utilization.

When IBM Storage Protect servers share a library, one server is set up as the library manager and controls library operations such as mount and dismount. The library manager also controls volume ownership and the library inventory. Other servers are set up as library clients and use server-to-server communications to contact the library manager and request resources.

Library clients must be at the same or an earlier version than the library manager server. A library manager cannot support library clients that are at a later version. For more information, see [Storage-agent and library-client compatibility with an IBM Storage Protect server](#).

LAN-free data movement

IBM Storage Protect provides the capability for a client, through a storage agent, to directly back up and restore data to a tape library on a SAN. This type of data movement is also known as LAN-free data movement.

Restriction: Centera storage devices cannot be targets for LAN-free operations.

[Figure 2 on page 18](#) shows a SAN configuration in which a client directly accesses a tape to read or write data.




Figure 2. LAN-free data movement

LAN-free data movement requires the installation of a storage agent on the client system. The server maintains the database and recovery log, and acts as the library manager to control device operations. The storage agent on the client handles the data transfer to the device on the SAN. This implementation frees up bandwidth on the LAN that would otherwise be used for client data movement.

Mixed device types in libraries

IBM Storage Protect supports mixing different device types within a single automated library, if the library can distinguish among the different media for the different device types. To simplify the configuration process, do not plan to mix different device types within a library. If you must mix device types, review the restrictions.

Libraries with this capability are models that have built-in mixed drives, or that support the addition of mixed drives. For more information, see the manufacturer's documentation. To learn about libraries that were tested on IBM Storage Protect with mixed device types, see the information for your operating system:

- [IBM Storage Protect Supported Devices for AIX, HP-UX, Solaris, and Windows](#)
- [IBM Storage Protect Supported Devices for Linux](#)

For example, you can have LTO Ultrium drives and IBM TS1100 drives in a single library that is defined to the IBM Storage Protect server.

Different media generations in a library

The IBM Storage Protect server allows mixed device types in an automated library, but the mixing of different generations of the same type of drive is generally not supported. New drives cannot write to the older media formats, and old drives cannot read new formats. LTO Ultrium drives and IBM 3592 tape drives are an exception to this rule.

If the new drive technology cannot write to media that is formatted by older generation drives, the older media must be marked read-only to avoid problems for server operations. Also, the older drives must be removed from the library, or the definitions of the older drives must be removed from the server. For example, the IBM Storage Protect server does not support the use of Oracle StorageTek 9940A drives with 9940B drives in combination with other device types in a single library.

In general, IBM Storage Protect does not support mixing generations of drives and media. However, the following mixtures of drives and media are supported:

- For the LTO Ultrium drives:
 - LTO Ultrium Generation 3 (LTO-3) with LTO Ultrium Generation 4 (LTO-4)
 - LTO Ultrium Generation 4 (LTO-4) with LTO Ultrium Generation 5 (LTO-5)
 - LTO Ultrium Generation 5 (LTO-5) with LTO Ultrium Generation 6 (LTO-6)
 - LTO Ultrium Generation 6 (LTO-6) with LTO Ultrium Generation 7 (LTO-7)
 - LTO Ultrium Generation 7 (LTO-7) media with LTO Ultrium Generation 8 (LTO-8 and LTO-M8) media in a library with LTO-8 tape drives or a library with mixed LTO-8 and LTO-7 tape drives
- For the IBM 3592 tape drives:
 - 3592 Generation 4 with 3592 Generation 3, for JB/JX or JA/JW media
 - 3592 Generation 5 with 3592 Generation 4, for JC/JY media
 - 3592 Generation 5A with 3592 Generation 5
 - 3592 Generation 6 with 3592 Generation 5A and/or 3592 Generation 5, for JD/JZ or JC/JY media

For detailed information, see the *IBM Tape Drive 3592 SCSI Reference* document.

The server supports these mixtures because the different drives can read and write to the different media. If you plan to upgrade all drives to the latest Generation, you must delete all existing drive definitions and the paths that are associated with them. Then, you can define the new Generation drives and paths.

Restrictions that apply to mixing LTO Ultrium tape drives and media

- LTO-5 drives can read only LTO-3 media. If you are mixing LTO-3 with LTO-5 drives and media in a single library, you must mark the LTO-3 media as read-only. You must check out all LTO-3 scratch volumes.
- LTO-6 drives can read only LTO-4 media. If you are mixing LTO-4 with LTO-6 drives and media in a single library, you must mark the LTO-4 media as read-only. You must check out all LTO-4 scratch volumes.
- LTO-7 drives can read only LTO-5 media. If you are mixing LTO-5 with LTO-7 drives and media in a single library, you must mark the LTO-5 media as read-only. You must check out all LTO-5 scratch volumes.
- LTO-8 drives are unable to read LTO-6 media. If you are mixing LTO-6 and LTO-8 drives and media in a single library, you must partition the library into two libraries. One library has only LTO-8 drives and media and the other has LTO-6 drives and media.

Restrictions that apply to mixed generation LTO Ultrium tape drives in a library

You must use tape cartridges that are an earlier generation than the tape drive. A later generation tape drive can read and write data to an earlier generation tape cartridge. For an example, if a library has LTO-7 and LTO-6 tape drives, you must use LTO-6 tape cartridges. Both the LTO-7 and LTO-6 tape drives can read and write data to LTO-6 tape cartridges.

Restrictions that apply to mixed generation LTO Ultrium tape cartridges in a library

You must use a tape cartridge that is the same generation as the tape drive, or one generation earlier. For example, if a library has LTO-7 tape drives, you can use LTO-7 tape cartridges or mixed LTO-7 and LTO-6 tape cartridges. If this library has LTO-7, LTO-6, and LTO-5 tape cartridges, you must change the access mode to READONLY for the LTO-5 tape cartridges.

To learn about additional considerations when you mix LTO Ultrium generations, see [“Defining LTO device classes”](#) on page 94.

When you use IBM Storage Protect, you cannot mix drives that are 3592, TS1130, TS1140, TS1150, and later drive generations. Use one of three special configurations. For details, see [“Defining 3592 device classes”](#) on page 97.

If you plan to encrypt volumes in a library, do not mix media generations in the library.

Mixed media and storage pools

You can optimize the efficiency of your tape solution by not mixing media formats in a storage pool. Instead of mixing formats, map each unique media format to a separate storage pool by using its own device class. This restriction also applies to LTO formats.

Multiple storage pools and their device classes of different types can point to the same library that can support them as described in [“Different media generations in a library”](#) on page 19.

You can migrate to a new generation of a media type within the same storage pool by following these steps:

1. Replace all older drives with the newer generation drives within the library. The drives should be mixed.
2. Mark the existing volumes with the older formats read-only if the new drive cannot append those tapes in the old format. If the new drive can write to the existing media in their old format, this is not necessary, but Step 1 is still required. If it is necessary to keep different drive generations that are read but not write compatible within the same library, use separate storage pools for each.

Definitions for Tape Storage Devices

Before the IBM Storage Protect server can use a tape device, you must configure the device to the operating system and to the server. As part of the planning process, determine the definitions for your Tape Storage Devices.

Tip: You can use the **PERFORM LIBACTION** command to simplify the process when you add devices to SCSI and VTL library types.

Table 7 on page 20 summarizes the definitions for different device types.

Table 7. Definitions for storage devices					
Device	Device types	Definitions			
		Library	Drive	Path	Device class
Magnetic disk	DISK	—	—	—	Yes ¹
	FILE ²	—	—	—	Yes
	<div><div>AIX</div><div>Windows</div>CENTERA</div> <div><div>Linux</div>CENTERA ³</div>	—	—	—	Yes

Table 7. Definitions for storage devices (continued)					
Device	Device types	Definitions			
		Library	Drive	Path	Device class
Tape	3590 3592 DLT LTO NAS VOLSAFE <div>AIX Windows</div> GENERICTAPE ECARTRIDGE ⁴	Yes	Yes	Yes	Yes
Removable media (file system)	REMOVABLEFILE	Yes	Yes	Yes	Yes

1. The DISK device class exists at installation and cannot be changed.
2. FILE libraries, drives, and paths are required for sharing with storage agents.
3.

Linux

 The CENTERA device type is only available for Linux x86_64 systems.
4. The ECARTRIDGE device type is for Oracle StorageTek cartridge tape drives such as 9840 and T10000 drives.

Planning the storage pool hierarchy

Plan the storage pool hierarchy to ensure that data is migrated daily from disk to tape. The migration releases space on the disk device and moves the data to tape for long-term retention. In this way, you can take advantage of the scalability, cost efficiency, and security features of tape storage.

Before you begin

The storage pool hierarchy helps to manage the flow of data. To understand the data flow, review [Figure 3](#) on [page 22](#).

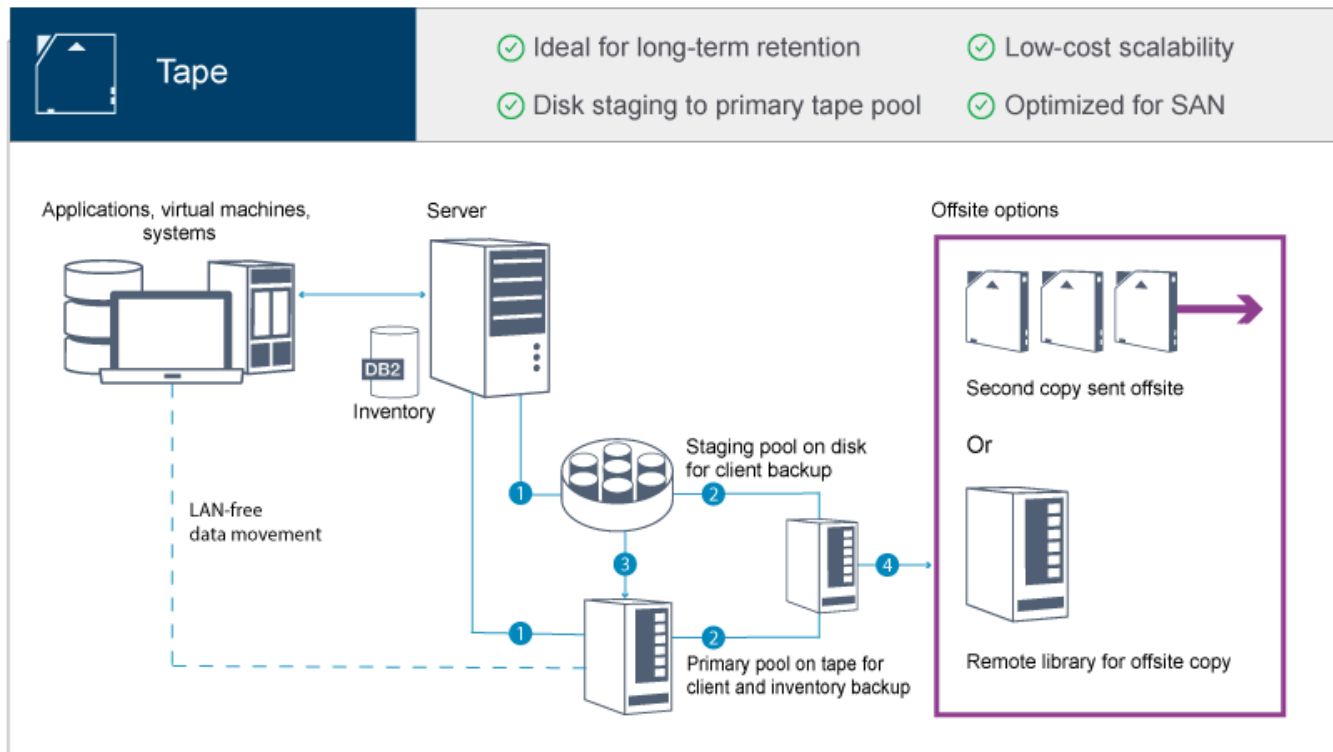


Figure 3. Tape solution

The following steps correspond to the numbers in the figure:

1. The server receives data from clients (applications, virtual machines, or systems) and stores the data on primary storage pools. Depending on the client type, the data is stored on a primary storage pool on disk or tape.
2. The data on disk and tape is backed up to a copy storage pool on tape.
3. Data in the primary storage pool on disk is migrated daily to the primary storage pool on tape.
4. Data from the copy storage pool on tape is moved offsite to support long-term retention and disaster recovery.

Procedure

To plan the storage pool hierarchy, answer the following questions:

- a. Which clients should back up data to disk, and which clients should back up data to tape?
 - The preferred method is to back up clients that host large objects, such as databases, to tape.
 - The preferred method is to back up all other clients to disk.
 - Virtual machine (VM) clients can be backed up to disk or tape. The preferred method is to back up a VM client to a separate disk storage pool, which is not migrated to tape. If you must migrate a VM client to tape, create a smaller disk storage pool to hold the VMware control files. This smaller disk storage pool cannot be allowed to migrate to tape. For more information about backing up a VM client to tape, see [Tape media guidelines](#) and [IBM Tivoli Storage Manager guest support for Virtual Machines and Virtualization](#).

Tip: If many clients must back up data to a single storage pool, consider using a storage pool on disk because you can specify many mount points. You can specify a maximum value of 999 for the **MAXNUMP** parameter on the **REGISTER NODE** command.

- b. What are the considerations for specifying the capacity of disk-based storage pools?

At minimum, plan enough capacity to store data from a single day of backup operations. The preferred method is to plan enough capacity to store data from two days' worth of backup operations and add a 20% buffer.

- c. What are the considerations for specifying the device class for the disk-based storage pool?

The preferred method is to specify a **FILE** device class. Set the **MOUNTLIMIT** parameter to 4000. Also, ensure that the node has a sufficiently high number of mount points, which you can specify by using the **MAXNUMMP** parameter on the **REGISTER NODE** command.

- d. Should data deduplication be specified for the disk storage pool?

No, because the data is stored on disk for only one day before the data is migrated to tape.

- e. Should automatic migration of data be specified based on a migration threshold?

No. Instead, plan to schedule daily migration by using the **MIGRATE STGPOOL** command. (To prevent automatic migration based on the migration threshold, specify a value of 100 for the **HIGHMIG** parameter and 0 for the **LOWMIG** parameter when you issue the **DEFINE STGPOOL** command.)

- f. Should a migration delay be specified?

The preferred method is to specify migration from disk to tape daily, and not specify a migration delay, which requires additional planning. For more information about migration delays, see [Migrating files in a storage pool hierarchy](#).

- g. How can the number of tape drives be calculated?

- i) Determine the native data transfer rate of the drive by reviewing the manufacturer's documentation. To obtain an estimate of the sustained data transfer rate in your storage environment, subtract 30% from the native data transfer rate.
- ii) Calculate the required rate of data ingestion by the server. Then, divide that figure by the sustained data transfer rate of a single tape device. The result is the minimum number of drives to support data ingestion.
- iii) Calculate the number of mount points that are required by clients that back up data to tape, including those clients that use multiple sessions. You can distribute the mount points over the backup window, keeping in mind that clients are likely backing up large objects, which might use most of the window.
- iv) Calculate the performance requirements *and* mount points that are required for maintenance tasks, such as disk-to-tape migration and tape-to-tape copies. By backing up data to tape, you can avoid migration processing, but making tape-to-tape copies will double the tape drive requirement.
- v) Calculate the number of additional drives that might be required, for example:
 - If a tape drive malfunctions, the issue impacts the number of available mount points and the ingestion rate. Consider provisioning spare drives. For example, if you require five tape drives for normal operations, consider provisioning two spare drives.
 - Restore and retrieve operations might require additional tape drives if you plan to run the operations simultaneously with data ingestion and maintenance operations. If necessary, provision additional tape drives and ensure that they are unused when you start the restore or retrieve operations.

- h. What alternatives are available for optimizing restore operations?

You can use collocation to improve system performance and optimize data organization. Collocation can reduce the number of volumes that must be accessed when a large amount of data must be restored:

- For disk-based storage pools, the preferred method is to use collocation by node. The server stores the data for the node on as few volumes as possible.
- For tape-based storage pools, the preferred method is to use collocation by group. Collocation by group results in a reduction of unused tape capacity, which allows for more collocated data on individual tapes.

For more information about collocation, see [“Optimizing operations by enabling collocation of client files” on page 164](#).

If you are an experienced system administrator, you might plan additional actions to optimize restore operations. See [Optimizing restore operations for clients](#), [File backup techniques](#), and [MOVE NODEDATA \(Move data by node in a sequential access storage pool\)](#).

Offsite data storage

To facilitate data recovery and as part of your disaster recovery strategy, store tape copies offsite.

Use the disaster recovery manager (DRM) function to configure and automatically generate a disaster recovery plan that contains the information, scripts, and procedures that are required to automatically restore the server and recover client data after a disaster. Choose from one of the following offsite data storage options as a disaster recovery strategy to protect tape copies:

Offsite vaulting from a single production site

Storage volumes, such as tape cartridges and media volumes, are vaulted at an offsite location. A courier transports the data from the offsite storage facility to the recovery site. If a disaster occurs, the volumes are sent back to the production site after hardware and the IBM Storage Protect server are restored.

Offsite vaulting with a recovery site

A courier moves storage volumes from the production site to an offsite storage facility. By having a dedicated recovery site, you can reduce recovery time compared to the single production site. However, this option increases the cost of disaster recovery because more hardware and software must be maintained. For example, the recovery site must have compatible tape devices and IBM Storage Protect server software. Before the production site can be recovered, the hardware and software at the recovery site must be set up and running.

Electronic vaulting

To use electronic vaulting as a disaster recovery strategy, the recovery site must have a running IBM Storage Protect server. Critical data is vaulted electronically from the production site to the recovery site. DRM is also used for offsite vaulting of noncritical data. Electronic vaulting moves critical data offsite faster and more frequently than traditional courier methods. Recovery time is reduced because critical data is already stored at the recovery site. However, because the recovery site runs continuously, the cost of the disaster recovery strategy is more expensive than offsite vaulting.

Related concepts

[Preparing for and recovering from a disaster by using DRM](#)

IBM Storage Protect provides a disaster recovery manager (DRM) function to recover your server and client data during a disaster.

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.

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 8 on page 25](#) provide examples about why you might want to assign varying levels of authority so that administrators can perform tasks:

Table 8. Scenarios for administrator roles	
Scenario	Type of administrator ID to set up
An administrator at a small company manages the server and is responsible for all server activities.	<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

Related tasks

[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.

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 tasks

[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.

Planning for storage of encrypted data

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

<i>Table 9. Selecting a data encryption method</i>		
Business requirement	Encryption method	Additional information
Protect data at the client level.	IBM Storage Protect client encryption	You can encrypt data at the file level by using an include/exclude list. In this way, you can maintain a high degree of control over which data is encrypted. Extra computing resources are required at the client that might affect the performance of backup and restore processes. For more information about this method, see client encryption .
Protect data in storage pool volumes on a tape drive.	Application method	When you use the Application method, IBM Storage Protect manages the encryption keys to protect data in storage pool volumes. You must take extra care to secure database backups because the encryption keys are stored in the server database. Without access to database backups and matching encryption keys, you cannot restore your data. You cannot use this method to encrypt database backups, exported data, or backup sets. For more information about the Application method, see “Tape encryption methods” on page 120 .
Protect data on a tape drive.	Library method	When you use the Library method, the library manages encryption keys. You can encrypt both data in storage pools and other data on a tape drive. You can control which volumes are encrypted by using their bar code serial numbers. For more information about the Library method, see “Tape encryption methods” on page 120 .
Protect data on a tape drive.	System method	When you use the System method, a device driver or the AIX operating system manages encryption. This encryption method is available only on the AIX operating system. You can encrypt both data in storage pools and other data on a tape drive. For more information about the System method, see “Tape encryption methods” on page 120 .

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 10 on page 27 describes the ports that are used by the server, client, and Operations Center.

Table 10. 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. The TCP PORT option listens for both TCP/IP and SSL-enabled sessions from the client. You can use the TCPADMINPORT option and ADMINONCLIENTPORT option to set port values for administrative client traffic.
RESTHTTPSPORT	8443	Inbound/ outbound	This port is used for secure communications between the hub server and the Operations Center.
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. A server can support both SSL and non-SSL communication by using the TCP PORT or TCPADMINPORT 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.
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 are specified for the source server on the DEFINE SERVER command.</p>
Client schedule port	Client port: 1501	Outbound	The client listens on the port that is named and communicates the port number to the server. The server contacts the client if server prompted scheduling is used. You can specify an alternative port number in the client options file.
Long-running sessions	KEEPALIVE setting: YES	Outbound	When the KEEPALIVE option is enabled, keepalive packets are sent during client/server sessions to prevent the firewall software from closing long-running, inactive connections.
Operations Center	HTTPS: 11090	Inbound	These ports are used for the Operations Center web browser. You can specify an alternative port number.

Table 10. Ports that are used by the server, client, and Operations Center (continued)

Item	Default	Direction	Description
Client management service port	Client port: 9028	Inbound	If you plan to use IBM Storage Protect client management services, 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.

Related information

[Collecting diagnostic information with IBM Storage Protect client management services](#)

[ADMINONCLIENTPORT server option](#)


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

[TCPADMINPORT server option](#)

[TCPPOINT server option](#)

Part 2. Implementation of a tape-based data protection solution

Implement the tape-based solution, which uses disk-to-disk-to-tape backup and disk staging to optimize storage. By implementing the tape solution, you can enable long-term data retention and achieve low-cost scalability.



Tips:

- In the described solution, data is *migrated* from disk storage pools to tape storage pools. However, instead of migrating the data, you can use the tiering-to-tape feature that was introduced in IBM Storage Protect Version 8.1.8. With this feature, you can automatically tier data from directory-container storage pools on disk to tape storage. You can specify that all data is tiered based on a specified age threshold, or that only inactive data is tiered based on an age threshold. For more information about tiering data to tape storage, see [Tiering data to cloud, tape, or file storage](#).
- The described solution does not include node replication. If you want to use node replication to back up a storage pool from disk to disk, ensure that the replication operation is completed before data is migrated from disk to tape. You can also use node replication to back up a storage pool on a local tape device to a copy storage pool on a local tape device.

Implementation roadmap

The following steps are required to set up a tape-based solution.

1. [Set up the system.](#)
2. [Install the server and the Operations Center.](#)
3. [Configure the server and the Operations Center.](#)
4. [Attach tape devices for the server.](#)
5. [Configure tape libraries for use by the server.](#)
6. [Set up a storage pool hierarchy.](#)

7. [Install and configure clients.](#)
8. [Configure LAN-free data movement.](#)
9. [Select an encryption method and configure encryption.](#)
10. [Set up tape storage operations.](#)
11. [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.

About this task

Tip: Procedures for setting up the server and the disk storage system are described. To get started with setting up tape devices, see [“Attaching tape devices for the server” on page 75.](#)

Configuring the storage hardware

To optimize disk storage, review the guidelines for setting up disk storage with IBM Storage Protect. Then, provide a connection between the server and the disk storage devices and complete other configuration tasks.

Before you begin

For guidelines about setting up disk storage, see [Checklist for storage pools on DISK or FILE](#)

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.
For more information, see [“Planning for disk storage” on page 13.](#)
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.
5. Ensure that tape and disk devices use different Host Bus Adapter (HBA) ports. Control tape and disk I/O by using the SAN. Use separate Fibre Channel ports for tape and disk I/O.

Related tasks

[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.

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
```

- 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

- Verify that user process resource limits, also known as *ulimits*, are set according to guidelines in [Table 11 on page 32](#). If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 11. 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 12 on page 35](#) 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 12. 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).

Table 12. Linux kernel parameter optimum settings (continued)	
Parameter	Description
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 13 on page 37](#). If ulimit values are not set correctly, you might experience server instability or a failure of the server to respond.

Table 13. 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

Complete the following steps to enable and configure multipathing for disk storage.

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
3321360050763008101057800000000000003004214503IBMfc
```

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

Complete the following steps to enable and configure multipathing for disk storage.

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

a) If needed, correct disk LUN host assignments and force a bus rescan.

For example:

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

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

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

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

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

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

Windows systems

Complete the following steps to enable and configure multipathing for disk storage.

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 tsmsrvrs -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>	
The database directories	<code>mkdir /tsminst1/</code> <code>TSMdbspace00</code> <code>mkdir /tsminst1/</code> <code>TSMdbspace01</code> <code>mkdir /tsminst1/</code> <code>TSMdbspace02</code> <code>mkdir /tsminst1/</code> <code>TSMdbspace03</code>	
Active log directory	<code>mkdir /tsminst1/TSMalog</code>	
Archive log directory	<code>mkdir /tsminst1/TSMarchlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir /tsminst1/</code> <code>TSMlogmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir /tsminst1/</code> <code>TSMarchlogfailover</code>	

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)	<code>mkdir C:\tsminst1</code>	
The database directories	<code>mkdir</code> <code>C:\tsminst1\TSMdbspace00</code> <code>mkdir</code> <code>C:\tsminst1\TSMdbspace01</code> <code>mkdir</code> <code>C:\tsminst1\TSMdbspace02</code> <code>mkdir</code> <code>C:\tsminst1\TSMdbspace03</code>	
Active log directory	<code>mkdir \tsminst1\TSMalog</code>	
Archive log directory	<code>mkdir C:\tsminst1\TSMarchlog</code>	
Optional: Directory for the log mirror for the active log	<code>mkdir</code> <code>C:\tsminst1\TSMlogmirror</code>	
Optional: Secondary archive log directory (failover location for archive log)	<code>mkdir</code> <code>C:\tsminst1\TSMarchlogfailove</code> <code>r</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 **mklv** command. The volume size, volume group, and device names vary, depending on the size of your system and variations in your disk configuration.

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

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

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

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

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

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

```
mount -a
```

7. List all file systems by issuing the **df** command.

Verify that file systems are mounted at the correct LUN and correct mount point. Also, verify the available space.

The 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 that you created in [“Creating the user ID and directories for the server instance” on page 41](#) has read and write access to the directories for the 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/36005076802810c50980000000000012
```

For storage pool LUNs, format xfs file systems:

```
mkfs -t xfs /dev/mapper/3600507630081010578000000000002c3
```

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 that you created in [“Creating the user ID and directories for the server instance” on page 41](#) has read and write access to the directories for the IBM Storage Protect server.

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 same 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 47 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.

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 same 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.

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 8 during IBM Storage Protect system setup 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 `dsmserve.opt` file.
2. Review the values in the following table and verify your server option settings, based on system size.

Server option	Value
ACTIVELOGDIRECTORY	Directory path that was specified during configuration
ACTIVELOGSIZE	131072
ARCHLOGCOMPRESS	No
ARCHLOGDIRECTORY	Directory path that was specified during configuration
COMMMETHOD	TCPIP
COMMTIMEOUT	3600
DEVCONFIG	<code>devconf.dat</code>
EXPINTERVAL	0

Server option	Value
IDLETIMEOUT	60
MAXSESSIONS	500
NUMOPENVOLSALLOWED	20
TCPADMINPORT	1500
TCPPORT	1500
VOLUMEHISTORY	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. To configure secure communications for the server, clients, and the Operations Center, verify the options in the following table.

Server option	All system sizes
SSLDISABLELEGACYTLS	YES
SSLFIPSMODE	NO
SSLTCPPORT	Specify the SSL port number. The server TCP/IP communication driver waits for requests on this port for SSL-enabled sessions from the client.
SSLTCPADMINPORT	Specify the port address on which the server waits for requests for SSL-enabled sessions from the command-line administrative client.
TLS12	YES

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

- Remove the asterisk at the beginning of a line to enable an option.
- 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.

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 14. 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.

Table 14. Password authentication characteristics (continued)	
Characteristic	More information
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 15. 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 **DSMADM** 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](#)

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

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 154.

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 TCPSPORT server option.

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.

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**.

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 figure provides an example of how to plan maintenance operations.




Figure 4. Daily schedule of server operations for a tape solution

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

Operation	Schedule
Client backup	Starts at 11 PM.
Storage pool backup	Starts at 6 AM.

Operation	Schedule
Processing for database and disaster recovery files	<ul style="list-style-type: none"> The database backup operation starts at 10 AM, or 11 hours after the beginning of the client backup operation. This process runs until completion. Device configuration information and volume history backup operations start at 5 PM, or 7 hours after the start of the database backup operation. Volume history deletion starts at 8 PM, or 10 hours after the start of the database backup operation.
Preparation of files for offsite vaulting	Starts at 10 AM, at the same time as processing for the database and disaster recovery files.
Migration from disk to tape	Starts at 12 PM, or 2 hours after the start of the database backup operation.
Inventory expiration	Starts at 2 PM, or 15 hours after the beginning of the client backup operation. This process runs until completion.
Space reclamation	Starts at 3 PM, or 16 hours after the beginning of the client backup operation.

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 operation before you create the schedule for database backups.

Use the **DEFINE DEVCLASS** command to create a device class that is named LTOTAPE:

```
define devclass ltotape devtype=lto library=ltolib
```

2. Set the device class for automatic database backups. Use the **SET DBRECOVERY** command to specify the device class that you created for the database backup in the preceding step. For example, if the device class is LTOTAPE, issue the following command:

```
set dbrecovery ltotape
```

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.

Operation	Example commands and additional information
Back up storage pools.	<p>Create a schedule to run the BACKUP STGPOOL command.</p> <p>For example, issue the following command to create a backup schedule for a primary storage pool that is named PRIMARY_POOL. The pool will be backed up to a copy storage pool, COPYSTG:</p> <pre>define schedule BACKUPSTGPOOL type=administrative cmd="backup stgpool primary_pool copystg" active=yes starttime=06:00 period=1</pre>

Operation	Example commands and additional information
Back up the database.	<p>Create a schedule to run the BACKUP DB command.</p> <p>For example, 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=ltotape type=full numstreams=3 wait=yes compress=yes" active=yes desc="Back up the database." startdate=today starttime=10:00:00 duration=45 durunits=minutes</pre>
Replicate nodes.	<p>Optionally, use node replication to protect client data by backing the data up to a failover server. For instructions, see Replicating client data to another server. Ensure that node replication is completed before migration operations begin.</p>
Migrate data from disk to tape daily.	<p>Create a schedule for storage pool migration.</p> <p>For example, if a disk storage pool is named DISKPOOL and the next storage pool is TAPEPOOL, you can schedule storage pool migration by issuing the following command:</p> <pre>define schedule stgpool_migration type=administrative cmd="migrate stgpool diskpool lomig=0" active=yes description="migrate disk storagepool to tapepool" startdate=today starttime=12:00 duration=2 durunits=hours period=1 perunits=days</pre> <p>To maximize throughput, you can specify the number of parallel processes to use for migrating files by completing the following steps:</p> <ol style="list-style-type: none"> For the tape storage pool, ensure that collocation is enabled. To verify whether collocation is enabled, run the QUERY STGPOOL command. Verify that a value of GROUP, NODE, or FILESPACE is specified in the COLLOCATE field. If a value of GROUP, NODE, or FILESPACE is not specified, use the UPDATE STGPOOL command to specify COLLOCATE=GROUP, COLLOCATE=NODE, or COLLOCATE=FILESPACE, depending on your system configuration. For the disk storage pool, use the DEFINE STGPOOL or UPDATE STGPOOL command to specify a value for the MIGPROCESS parameter. For example, if you have 12 tape drives, specify MIGPROCESS=10. In this way, a maximum of 10 tape drives are used for migration processes. Two drives are reserved for other tasks, such as restore, database backup, and client backup operations.
Prepare files for offsite vaulting.	<ol style="list-style-type: none"> Move tape volumes offsite by following the instructions in “Moving backup media” on page 62. Create the disaster recovery plan file by issuing the PREPARE command on the source server: <pre>prepare</pre> Ensure that all volumes that are required for disaster recovery are included in the recovery plan file. For more information, see “Preparing for and recovering from a disaster by using DRM” on page 208.

Operation	Example commands and additional information
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 to be equal to the number of processor cores that you specified for your system.</p> <p>For example, 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=8 duration=120" active=yes desc="Remove expired objects." startdate=today starttime=14:00:00 duration=1 durunits=hours</pre>
Reclaim space.	<p>Create a schedule to run the RECLAIM STGPOOL command.</p> <p>For example, issue the following command to create a schedule that is named RECLAIM:</p> <pre>define schedule RECLAIM type=admin cmd="reclaim stgpool tapepool duration=60" startdate=today starttime=15:00:00 duration=5 durunits=hours</pre> <p>Tip: To maximize throughput, you can specify the number of parallel processes to use for reclaiming space. Update the tape storage pool by using the UPDATE STGPOOL command and specify a value for the RECLAIMPROCESS parameter. For example, if you have 12 tape drives, specify RECLAIMPROCESS=5. Because two drives are used for each reclamation process, the total number of drives that can be used for reclamation is 10. Two drives are reserved for backup operations.</p>

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

[UPDATE STGPOOL](#) (Update a storage pool)

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

[DEFINE STGPOOL](#) (Define a volume in a storage pool)

Moving backup media

To recover from a disaster, you need database backup volumes, copy storage pool volumes, and additional files. To stay prepared for a disaster, you must complete daily tasks.

Before you begin

To display all virtual copy storage pool and database backup volumes that have their backup objects on the remote target server, issue the **QUERY DRMEDIA** command:

```
query drmedia * wherestate=remote
```

About this task

The disaster recovery manager (DRM) function enables you to track the movement of offsite media. The following figure shows the lifecycle of a typical operation to move backup media offsite and back onsite, as part of disaster recovery operations.

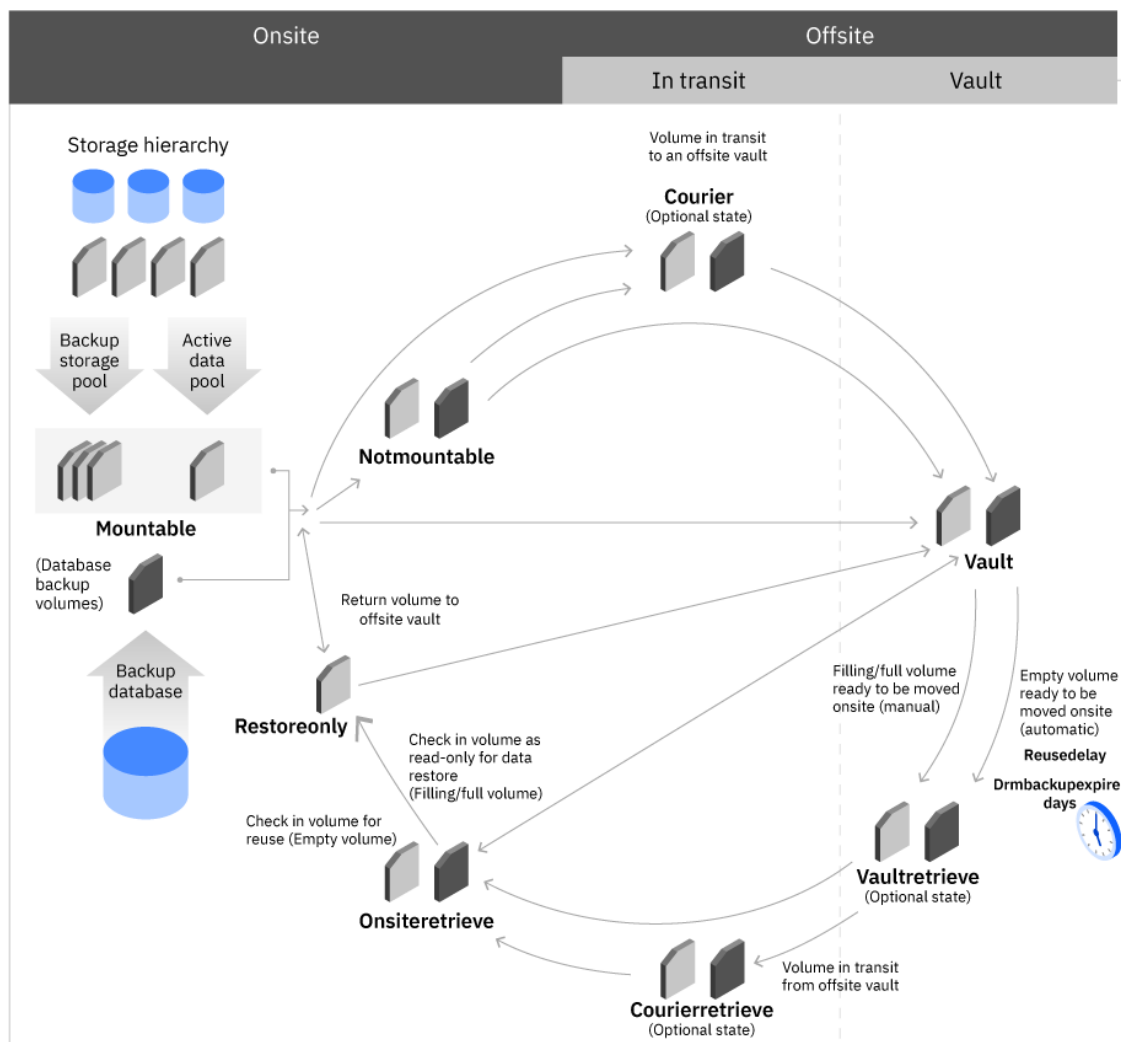


Figure 5. Offsite and onsite movement of backup volumes

DRM assigns the following media states to volumes. The media states track a volume as it moves from location to location. Some media states are optional. Depending on how finely your organization wants to track a volume's movements, your organization might skip these optional media states. The following media states are shown:

MOUNTABLE

The volume contains valid data, is onsite, and can be accessed by the IBM Storage Protect server.

NOTMOUNTABLE

The volume contains valid data, is onsite, but cannot be accessed by the IBM Storage Protect server.

COURIER

The volume contains valid data and is in transit to the vault.

VAULT

The volume contains valid data and is at the vault.

VAULTRETRIEVE

The volume, which is at the offsite vault, no longer contains valid data and is to be returned onsite.

COURIERRETRIEVE

The volume no longer contains valid data and is returned by the courier.

ONSITERETRIEVE

The volume no longer contains valid data and is moved back onsite. The volume records of database backup, scratch copy storage pool volumes, and scratch active-data pool volumes are deleted from the database. For private copy storage pool volumes and active-data pool volumes, the access mode is updated to READWRITE.

RESTOREONLY

The volume is checked into the library to enable restoration of data. The volume is used only for data restoration.

Moving copy storage pool volumes offsite

You can send your backup media offsite after you create the backup copies of your primary storage pools and database. To send media offsite, mark the volumes as unavailable to IBM Storage Protect and give them to the courier.

Before you begin

Ensure that storage pool backup processes are completed. In this way, you can avoid issues that might occur when the **MOVE DRMEDIA** and **BACKUP STGPOOL** commands are run concurrently.

Restriction: To move retention volumes offsite, that is, tape volumes that contain retention set data, do not use this procedure. You must use the **MOVE RETMEDIA** command or move media operations in the Operations Center. For instructions, see [“Moving retention set data to and from tape storage” on page 67](#).

Procedure

1. To identify the copy storage pool and database backup volumes to be moved offsite, issue the **QUERY DRMEDIA** command and specify the **WHERESTATE** parameter.

```
query drmedia * wherestate=mountable
```

2. Indicate the movement of volumes whose current state is MOUNTABLE by issuing the **MOVE DRMEDIA** command and specifying the **WHERESTATE** parameter:

```
move drmedia * wherestate=mountable
```

- a) During check-out processing, SCSI libraries request operator intervention. Bypass these requests and eject the cartridges from the library by issuing the following command:

```
move drmedia * wherestate=mountable remove=no
```

- b) Access a list of the volumes to identify and remove the cartridges from the library by issuing the following command:

```
query drmedia wherestate=notmountable
```

For all volumes in the MOUNTABLE state, DRM completes the following tasks:

- Updates the volume state to NOTMOUNTABLE and, if you issued the **SET DRMNOTMOUNTABLENAME** command, updates the volume location. If you do not issue the **SET DRMNOTMOUNTABLENAME** command, the default location is NOTMOUNTABLE.
- Updates the access mode to unavailable for a copy storage pool volume.
- Checks volumes out of automated libraries.

3. Send the volumes to the courier for transit to the offsite location and issue the following command:

```
move drmedia * wherestate=notmountable
```

For all volumes in the NOTMOUNTABLE state, DRM updates the volume state to COURIER and the volume location according to the **SET DRMCOURIERNAME** command. If you did not issue the **SET** command, the default location is COURIER.

Tip: You can avoid going through all volume states by issuing the **MOVE DRMEDIA** command and specifying the **TOSTATE** parameter setting to name the destination state. For example, to change volumes from the NOTMOUNTABLE state to the VAULT state, issue the following command:

```
move drmedia * wherestate=notmountable tostate=vault
```

For all volumes in the NOTMOUNTABLE state, DRM updates the volume state to VAULT and updates the volume location according to the **SET DRMVAULTNAME** command. If the **SET** command is not yet issued, the default location is VAULT.

4. When receipt of the volumes is confirmed at the vault location, issue the **MOVE DRMEDIA** command to specify the COURIER state:

```
move drmedia * wherestate=courier
```

For all volumes in the COURIER state, DRM updates the volume state to VAULT and the volume location according to the **SET DRMVAULTNAME** command. If you did not issue the **SET** command, the default location is VAULT.

5. Display a list of volumes that contain valid data at the vault by issuing the following command:

```
query drmedia wherestate=vault
```

Moving copy storage pool volumes onsite

You can move backup media onsite, as part of disaster recovery operations. You can return the volumes onsite to restore data. You can also expire non-virtual database backup volumes and return the volumes onsite for reuse or removal.

Before you begin

Restriction: To return volumes that contain set retention set data onsite, do not use this procedure. You must use the **MOVE RETMEDIA** command or move media operations in the Operations Center. For instructions, see [“Moving retention set data to and from tape storage” on page 67](#).

If you are returning volumes for reuse, confirm that expiration dates are reached for these volumes. You can expire a database backup volume when all of the following conditions are true:

- The age of the last volume of the series exceeds the expiration value. The expiration value is the number of days since the last backup in the series. At installation, the expiration value is 60 days. To override this value, you can issue the **SET DRMDBBACKUPEXPIREDAYS** command.
- All volumes in the series are in the VAULT state.

- The volume is not part of the most recent database backup series.

Start expiration processing manually by issuing the **EXPIRE INVENTORY** command or automatically by using the EXPINTERVAL option setting that is specified in the server options file.

Procedure

To move storage pool volumes onsite, take one of the following actions:

Task	Procedure
Move an empty volume onsite for reuse or removal.	<p>To move empty storage pool volumes onsite, complete the following steps:</p> <ol style="list-style-type: none"> Specify the number of days before a database backup series is expired by issuing the SET DRMDBBACKUPEXPIREDAYS command. For example, to set the number of days to 30, issue the following command: <pre>set drmdbbackupexpiredays 30</pre> <p>Tip: Issue the DEFINE STGPPOOL command and specify the same value for the REUSEDELAY parameter in your copy storage pool definition to ensure that the following occurs:</p> <ul style="list-style-type: none"> • The database can be restored to an earlier level. • Database references to files in the copy storage pool are still valid. <p>If copy storage pools that are managed by DRM have different REUSEDELAY values, issue the SET DRMDBBACKUPEXPIREDAYS command and set the REUSEDELAY parameter to the highest value.</p> Identify all volumes at the offsite vault that no longer contain valid data and can be returned to onsite. For empty volumes, the server automatically places the volume in the media state VAULTRETRIEVE. Issue the following command: <pre>query drmedia * wherestate=vaultretrieve</pre>

Task	Procedure
	<p>c. To start the process of moving a copy storage pool, issue the following command:</p> <pre>move drmedia * wherestate=vaultretrieve</pre> <p>Restriction: A copy storage pool volume can be moved onsite if it is in the EMPTY state for at least the number of days that is specified by the REUSEDelay parameter on the DEFINE STGPOOL command.</p> <p>The server completes the following actions for all volumes in the VAULTRETRIEVE state:</p> <ul style="list-style-type: none"> • Changes the volume state to COURIERRETRIEVE • Updates the location of the volume according to the value that is specified in the SET DRMCOURIERNAME command <p>Tip:</p> <p>You can also specify the destination of the volumes by issuing the MOVE DRMEDIA command and specifying the TOSTATE parameter setting. For example, to move volumes from the VAULTRETRIEVE state to the ONSITERETRIEVE state, issue the following command:</p> <pre>move drmedia * wherestate=vaultretrieve tostate=onsiteretrieve</pre> <p>The server completes the following actions for all volumes in the VAULTRETRIEVE state:</p> <ul style="list-style-type: none"> • Moves the volumes onsite where they can be reused or removed • Deletes the database backup volumes from the volume history table • Deletes the record in the database for scratch copy storage pool volumes and, for private copy storage pool volumes, updates the access to read/write <p>d. When the courier returns the volumes onsite, issue the following command:</p> <pre>move drmedia * wherestate=courierretrieve</pre> <p>The server completes the following actions for all volumes in the COURIERRETRIEVE state:</p> <ul style="list-style-type: none"> • Moves the volumes onsite where they can be reused or disposed of • Deletes the database backup volumes from the volume history table • Deletes the record in the database for scratch copy storage pool volumes. For private copy storage pool volumes, updates the access to read/write
Move a non-empty volume onsite for data restoration.	<p>To move storage pool volumes onsite to restore data, complete the following steps:</p> <p>a. Identify the volumes that you want to return onsite. To locate a required volume at its offsite location, issue the QUERY DRMEDIA command and specify the WHERESTATE parameter. For example, to view all volumes that are located in the offsite vault, issue the following command:</p> <pre>query drmedia * wherestate=vault</pre>

Task	Procedure
	<p>b. Move the volume onsite. Specify the destination of the volume by issuing the MOVE DRMEDIA command with the TOSTATE parameter setting. For example, to move volume VOL001 onsite, issue the following command:</p> <pre>move drmedia vol001 wherestate=vault tostate=onsiteretrieve</pre> <p>The server completes the following actions for all specified volumes in the VAULT state:</p> <ul style="list-style-type: none"> • Moves the volumes onsite where they can be used to restore data • Changes the volume state to ONSITERETRIEVE <p>c. Check in the volume to the tape library and make the volume available for restore operations. To move the volume from the ONSITERETRIEVE state to the RESTOREONLY state, issue the CHECKIN LIBVOL command. For example, if the name of the library is LIBNAME, you could issue the following command:</p> <pre>checkin libvol libname search=bulk waittime=0 checklabel=barcode status=private</pre> <p>Tip: For tape volumes in SCSI libraries, you can decrease the check-in time by specifying that the server reads the bar code label.</p> <p>The volume is added to an automated library and the volume's media state changes to RESTOREONLY.</p> <p>d. Restore the data from the tape volume. After data restoration is completed, you can send the tape volumes back to the offsite vault again. You can process the tape volume with other tape volumes that are being moved offsite. The volume's media state changes to MOUNTABLE by default. Issue the following command:</p> <pre>move drmedia * wherestate=restoreonly</pre> <p>Alternatively, you can specify the destination of the volumes by issuing the MOVE DRMEDIA command and specifying the TOSTATE parameter setting. For example, to move volumes from the RESTOREONLY state to the VAULT state, issue the following command:</p> <pre>move drmedia * wherestate=restoreonly tostate=vault</pre>

Results

The selected storage pool volumes are returned onsite and checked into the tape library. Empty tape volumes are returned to scratch status and are available for reuse. Non-empty volumes are in the RESTOREONLY state and can be used to restore the data.

Moving retention set data to and from tape storage

You can copy retention set data to tape volumes, which you can move from an onsite library to an offsite tape storage vault. Vaults are designed to provide long-term, secure storage. After the retention set is copied to tape and the tape volume is removed from the tape library, you can track the movement of the volume offsite and onsite.

A tape volume that contains data for one or more retention sets is called a *retention volume*. As the tape volume is moved from one location to another, the volume state changes to reflect the new location and you can use this information to track the volume's physical location.

The lifecycle of a retention volume consists of the following main stages:

1. When the process to write a retention set to a tape volume begins, either a scratch volume is acquired from the tape library scratch pool or an existing volume is selected from the retention pool. Data from one or more retention sets is written to the volume. When the volume is full, it is taken by courier to an offsite vault.
2. If the volume contains data that must be restored, the volume is retrieved from the vault and brought back onsite by courier. After the data in the retention set is restored, the volume is moved back to the offsite vault.
3. Over time, data in retention sets can expire, based on expiration policies. If expiration dates are reached for all of the retention sets that have data on the volume, the volume can be returned onsite for reuse.

The following figure shows the lifecycle of a typical operation to move retention volumes offsite and back onsite.




Figure 6. Offsite and onsite movement of retention volumes

The volume's media state helps you identify its current location as it is moved from your onsite library to an offsite vault, and then back onsite for either data restoration or tape reuse. The volume's media state is a logical designation that is related to the volume's physical location. Some media states are optional.

Depending on how finely your organization wants to track a volume's movements, your organization might skip these optional media states. The following media states are shown:

MOUNTABLE

The volume is onsite and is checked into the library. Data from one or more retention sets is written to the volume.

NOTMOUNTABLE

The volume is onsite, but checked out of the library and ready to be sent offsite.

COURIER

The volume is in transit to an offsite vault.

VAULT

The volume is in an offsite vault for long-term storage.

VAULTRETRIEVE

The volume is ready to be moved back onsite from an offsite vault. Empty volumes can be brought onsite and reused. The server detects that the volume contains only expired data and automatically places the volume in the VAULTRETRIEVE media state. Filling or Full volumes can also be brought onsite for data restoration, but you must specify this action by using the **TOSTATE** parameter setting on the **MOVE RETMEDIA** command.

COURIERRETRIEVE

The volume is being moved back onsite from an offsite vault.

ONSITERETRIEVE

The volume was retrieved from the offsite vault and is back onsite. Non-empty volumes can be checked into the library to restore retention set data from the volume. Empty volumes can be checked in and reused.

RESTOREONLY

The volume is checked into the library to enable restoration of retention set data.

Moving retention volumes offsite

You can send retention volumes that contain data from one or more retention sets to an offsite location. Offsite vaults are designed to provide secure storage for tape volumes and help to ensure that the data can be restored if necessary.

Before you begin

Tip: If you do not use the **MOVE DRMEDIA** command to move database backup volumes offsite and onsite, you can also use the **MOVE RETMEDIA** command to do so. For more information, see [“Moving copy storage pool volumes offsite”](#) on page 63.

- After the retention set that you want to send offsite is created, back up the server database volumes by issuing the **BACKUP DB** command. If you want to ensure that the database backup volume is sent offsite along with the retention volume, you must specify the **SOURCE** parameter on the **MOVE RETMEDIA** command.

Restriction: You cannot use move media operations in the Operation Center to send a database backup volume offsite. Database backup volumes are moved by using the **MOVE RETMEDIA** command.

For information about using the Operations Center to move retention volumes, see the Operations Center online help.

- Ensure that the retention sets that you want to copy have a status of Completed. This status indicates that the retention sets are fully copied to tape and the tape volumes can be moved to an offsite vault. In this way, you can avoid issues that might occur when move media and copy retention set operations are run concurrently.

Procedure

1. Identify the retention storage pool and database backup volumes to be moved offsite by issuing the **QUERY RETMEDIA** command:

```
query retmedia * wherestate=mountable
```

2. Initiate the movement of volumes whose current state is MOUNTABLE. By default, all non-empty volumes are included whether they belong to retention sets that are being copied or to retention sets that are fully copied. Issue the following command:

```
move retmedia * wherestate=mountable
```

- a) If you are using a SCSI library, during check-out processing, SCSI libraries request operator intervention. Bypass these requests and eject the cartridges from the library by issuing the following command:

```
move retmedia * wherestate=mountable remove=no
```

- b) Get a list of the volumes to identify and remove from the library by issuing the following command:

```
query retmedia wherestate=notmountable
```

For all volumes in the MOUNTABLE state, the **MOVE RETMEDIA** command completes the following tasks:

- Updates the volume state to NOTMOUNTABLE and, if you issued the **SET DRMNOTMOUNTABLENAME** command, updates the volume location. If you did not issue the **SET DRMNOTMOUNTABLENAME** command, the default location is NOTMOUNTABLE.
- Updates the volume access mode to unavailable.
- Checks volumes out of automated libraries.

Tip: Depending on how finely your organization wants to track a volume's movements, your organization might skip some media states. You can avoid going through all the different media states by specifying the **TOSTATE** parameter on the **MOVE RETMEDIA** command to name the destination state. For example, to change the volumes directly from NOTMOUNTABLE state to VAULT state, issue the following command:

```
move retmedia * wherestate=notmountable tostate=vault
```

3. Send the volumes to the courier for transit to the offsite location and issue the following command:

```
move retmedia * wherestate=notmountable
```

For all volumes in the NOTMOUNTABLE state, the volume state is updated to the COURIER state and the volume location is updated according to the **SET DRMCOURIERNAME** command. If you did not issue the **SET DRMCOURIERNAME** command, the default location is COURIER.

4. Track the movement of the tape volume while it is in transit to an offsite vault. Issue the following command:

```
query retmedia * wherestate=courier
```

5. When the vault location confirms receipt of the volumes, issue the **MOVE RETMEDIA** command to specify the COURIER state:

```
move retmedia * wherestate=courier
```

For all volumes in the COURIER state, the volume state is updated to VAULT and the volume location is updated according to the **SET DRMVAULTNAME** command. If you did not issue the **SET DRMVAULTNAME** command, the default location is VAULT.

For all volumes in the NOTMOUNTABLE state, the **MOVE RETMEDIA** command updates the volume state to VAULT and the volume location is updated according to the **SET DRMVAULTNAME** command. If the **SET DRMVAULTNAME** command is not yet issued, the default location is VAULT.

Results

The retention volumes and any specified database backup volumes are moved to the offsite tape vault. If the retention set data must be restored, the volumes can be retrieved from the vault.

Moving retention volumes onsite

If a retention set must be restored, you can bring the tape volumes that contain the retention set data onsite for restore operations. If the expiration dates are reached for all of the retention sets that have data on the retention volume, you can bring back the empty volumeback onsite to be reused.

Before you begin

If you are returning empty volumes for reuse, confirm that expiration dates are reached for all of the retention sets that have data on the volume and the retention set is expired. You can start expiration processing manually by issuing the **EXPIRE INVENTORY** command or you can use the **DELETE RETSET** command to mark the retention set for deletion.

Tip: If you do not use the **MOVE DRMEDIA** command to move database backup volumes offsite and onsite, you can also use the **MOVE RETMEDIA** command to do so. For more information, see [“Moving copy storage pool volumes offsite”](#) on page 63.

Procedure

Complete the following steps to move retention volumes onsite.

Task	Procedure
Move an empty volume onsite for reuse.	<p>To move empty retention volumes onsite, complete the following steps:</p> <ol style="list-style-type: none"> Identify the retention volumes in the offsite vault that you want to return onsite. For empty volumes, the server detects that the volume contains only expired data and automatically places the volume in the media state VAULTRETRIEVE. Issue the following command: <pre>query retmedia * wherestate=vaultretrieve volstatus=empty</pre> Move the tape volumes onsite. Specify the destination of the volumes by issuing the MOVE RETMEDIA command and specifying the TOSTATE parameter. Issue the following command: <pre>move retmedia * wherestate=vaultretrieve volstatus=empty tostate=onsiteretrieve</pre> <p>Restriction: A retention storage pool volume can be moved onsite if it is in the EMPTY state for at least the number of days that are specified by the REUSEDELAY parameter on the DEFINE STGPOOL command.</p> <p>The server completes the following actions:</p> <ul style="list-style-type: none"> Changes the volume state to ONSITERETRIEVE Deletes the database backup volumes from the volume history table Deletes the record in the database for scratch retention volumes Check in the empty volume to the tape library and make available for reuse by issuing the CHECKIN LIBVOL command and specifying the volume as a scratch volume. <p>Tip: For tape volumes in SCSI libraries, you can decrease the check-in time by specifying that the server reads the bar code label.</p> <p>Issue the following command:</p> <pre>checkin libvol libname search=bulk waittime=0 checklabel=barcode status=scratch</pre>

Task	Procedure
<p>Move a non-empty volume onsite for data restoration.</p>	<p>To move retention volumes onsite for data restore, complete the following steps:</p> <ol style="list-style-type: none"> Identify the volumes that contain the retention set data that you want to restore. <ul style="list-style-type: none"> To identify the volumes that are used by each retention set, issue the following command: <pre data-bbox="617 430 987 462">query retset listvolumes=yes</pre> To identify the retention sets that have data on a retention volume, issue the following command: <pre data-bbox="617 577 987 609">query volume listretsets=yes</pre> Locate the required volume at its offsite location by issuing the QUERY RETMEDIA command and specifying the WHERESTATE parameter. For example, to view all volumes that are located in the offsite vault, issue the following command: <pre data-bbox="604 781 1029 812">query retmedia * wherestate=vault</pre> Move the required volume onsite. Specify the destination of the volumes by issuing the MOVE RETMEDIA command and specifying the TOSTATE parameter. For example, to move volume VOL001 onsite, issue the following command: <pre data-bbox="604 987 1370 1018">move retmedia VOL001 wherestate=vault tostate=onsiteretrieve</pre> <p>Important: To be eligible for reclamation processing, retention storage pool volumes must be in the MOUNTABLE state. Reclamation processing does not reclaim volumes that are in the ONSITERETRIEVE or RESTOREONLY states. If you return retention storage pool volumes onsite by issuing the MOVE RETMEDIA command and specifying either the TOSTATE=ONSITERETRIEVE or TOSTATE=RESTOREONLY parameter values, storage reclamation processing skips these volumes.</p> Check in the volume to the tape library and make the volume available for restore operations. To ensure that the volume can be used only for data restoration, its access mode is read only. To move the volume from the ONSITERETRIEVE state to the RESTOREONLY state, issue the CHECKIN LIBVOL command. Issue the following command: <pre data-bbox="604 1459 1107 1501">checkin libvol libname search=bulk waittime=0 checklabel=barcode status=private</pre> <p>Tip: For tape volumes in SCSI libraries, you can decrease the check-in time by specifying that the server reads the bar code label.</p> <p>The volume is added to an automated library and the volume's media state changes to RESTOREONLY.</p>

Results

The selected retention volumes are returned onsite and checked into the tape library. Empty tape volumes are returned to scratch status and are available for reuse. Non-empty volumes are in the RESTOREONLY state and can be used to restore the data.

What to do next

After data restoration is completed, you can send the tape volumes back to the offsite vault again. Issue the following command:

```
move retmedia * wherestate=restoreonly tostate=vault
```

Alert messages to monitor movement of retention volumes

If you send retention volumes offsite or back onsite, the IBM Storage Protect server generates alerts in the form of ANR messages to report any issues and to help you monitor the status.

To view all messages, see the IBM Storage Protect error log. For detailed documentation about messages, see ANR Messages. Commonly issued messages are described in the following table:

Table 16. Sending retention tape volumes to an offsite vault		
Action	ANR message	Description
The retention set is copied to the tape volume.	ANR3852I	This information message indicates that the retention set was successfully copied to the tape volume. Details of the copy operation are provided. The retention set state is COMPLETED.
Tape volumes are checked out of a tape library.	ANR6697I	This information message indicates that tape volumes in a MOUNTABLE state were successfully checked out of a tape library.
The tape volume is checked out of the library and moved from a MOUNTABLE state to a VAULT state.	ANR6683I	This information message indicates that retention data was successfully moved and the state was changed.

Table 17. Checking in tape volumes to the tape library to use for restore operations		
Action	ANR message	Description
A retention volume that contains data was checked in to the onsite tape library successfully.	ANR8532I	<p>This information message indicates that a volume with data is successfully checked into the onsite tape library. For retention volumes, the media state of the volume changes from ONSITERETRIEVE to RESTOREONLY and its access mode is read only. The retention set data in the volume can now be restored.</p> <p>Tip: This message does not appear if the tape volume that is being checked in is empty.</p>
You are attempting to check in a non-empty retention volume as a scratch volume to the tape library.	ANR8443E	This error message is triggered because a retention volume that contains data cannot be checked into a tape library and assigned a status of SCRATCH. The volume is not checked in and the data on the tape is not overwritten.

Table 18. Checking in expired retention volumes to a tape library

Action	ANR message	Description
An empty retention volume is checked in to an onsite tape library.	ANR8430I	This information message indicates that an empty volume is successfully checked into an onsite tape library. The volume is returned to scratch status.
An attempt to check in an empty retention volume to an onsite tape library failed.	ANR8832E	This error message indicates that an operation to check an empty retention volume in to an onsite tape library failed.

Defining client schedules

Use the Operations Center to create schedules for client operations.

Procedure

1. On the Operations Center menu bar, hover over **Clients**.
2. Click **Schedules**.
3. Click **+Schedule**.
4. Complete the steps in the **Create Schedule** wizard.

Set client backup schedules to start at 22:00, based on the server maintenance activities that you scheduled in [“Defining schedules for server maintenance activities”](#) on page 58.

Attaching tape devices for the server

Before the server can use a tape device, you must attach the device to your server system and install the appropriate tape device driver.

About this task

To optimize system performance, use fast, high-capacity tape devices. Provision enough tape drives to meet your business requirements.

Attach tape devices on their own host bus adapter (HBA), not shared with other devices types such as disk. IBM tape drives have some special requirements for HBAs and associated drivers.

Restrictions:

- IBM LTO-9 tape drives are supported. However, when a new tape volume is inserted into an LTO-9 tape drive, the drive must load and initialize the tape. Typically, the loading and initialization process can be completed in two hours, but the time can vary, based on several factors. The initialization process continues until the tape drive is ready. You cannot cancel the process. To avoid this issue, manually load the volumes into the tape drive in advance, before the tape drive is used by the server. To manually load the tape volumes, you can use the IBM Tape Diagnostic Tool (ITDT) or the IBM Storage Protect **lbtest** utility.
- IBM LTO-9 tape drives require the ULTRIUM9 recording format for uncompressed mode and the ULTRIUM9C recording format for compressed mode.

Attaching an automated library device to your system

You can attach an automated library device to your system to store your data on tapes.

About this task

Before you attach an automated library device, consider the following restrictions:

- Attached devices must be on their own Host Bus Adapter (HBA).
- An HBA must not be shared with other device types, such as a disk.
- For multiport Fibre Channel HBAs, devices must be attached on their own port. These ports must not be shared with other device types.
- IBM tape drives have some special requirements on HBA and associated drivers. For more information about devices, see the website for your operating system:
 - [IBM Storage Protect Supported Devices for AIX](#)
 - [IBM Storage Protect Supported Devices for Linux and Windows](#)

Procedure

To use the Fibre Channel (FC) adapter, complete the following steps:

1. Install the FC adapter and associated drivers.
2. Install the appropriate device drivers for attached medium changer devices.

Related concepts

[Selecting a tape device driver](#)

To use tape devices with IBM Storage Protect, you must install the appropriate tape device driver.

Setting the library mode

For the IBM Storage Protect server to access a SCSI library, the tape device must be set for the appropriate mode.

About this task

Some libraries have front panel menus and displays that can be used for explicit operator requests. However, if you set the tape device to respond to such requests, the device typically does not respond to IBM Storage Protect requests.

Some libraries can be placed in SEQUENTIAL mode, in which volumes are automatically mounted in drives by using a sequential approach. This mode conflicts with how IBM Storage Protect accesses the tape device. A library that is configured in SEQUENTIAL mode is not detected by the system device driver as a library changer device, IBM tape device driver, and IBM Storage Protect tape device driver.

Procedure

1. Refer to the documentation for your tape device to determine how to set the library mode.
2. Set the mode to the appropriate mode for your tape device. For most tape devices, the appropriate mode is called the RANDOM mode. If your tape device does not have a RANDOM mode, consult the documentation for your device to identify the appropriate mode.

Selecting a tape device driver

To use tape devices with IBM Storage Protect, you must install the appropriate tape device driver.

Related reference

[Installing and configuring tape device drivers](#)

Before you can use tape devices with IBM Storage Protect, you must install the correct tape device driver.

IBM tape device drivers

IBM tape device drivers are available for most IBM labeled tape devices.

You can download IBM tape device drivers from the Fix Central website:

1. Go to the Fix Central website: [Fix Central website](#).
2. Click **Select product**.
3. Select **System Storage** for the **Product Group** menu.
4. Select **Tape systems** for the **System Storage** menu.
5. Select **Tape drivers and software** for the **Tape systems** menu.
6. Select **Tape device drivers** for the **Tape drivers and software** menu. In addition to tape drivers, you also get access to tools such as the IBM Tape Diagnostic Tool (ITDT).
7. Select your operating system for the **Platform** menu.

AIX | Windows

For the most up-to-date list of devices and operating-system levels that are supported by IBM tape device drivers, see the IBM Storage Protect Supported Devices website at [Supported devices for AIX and Windows](#).

Linux

For the most up-to-date list of tape devices and operating-system levels that are supported by IBM tape device drivers, see the IBM Storage Protect Supported Devices website at [Supported devices for Linux](#).

IBM tape device drivers support only some Linux kernel levels. For more information about supported kernel levels, see the [Fix Central website](#).

IBM Storage Protect tape device drivers

The IBM Storage Protect server provides tape device drivers.

An IBM Storage Protect tape device driver is installed with the server.

AIX

You can use the generic SCSI tape device driver that is provided by the IBM AIX operating system to work with tape devices that are not supported by the IBM Storage Protect device driver. If the AIX generic SCSI tape device driver is used, the GENERICTAPE device class must be set to the device type that is specified in the **DEFINE DEVCLASS** command.

For the following tape devices, you can choose whether to install the IBM Storage Protect tape device driver or the native device driver for your operating system:

ECART
LTO (not from IBM)

All SCSI-attached libraries that contain tape drives from the list must use the IBM Storage Protect changer driver.

Tape device drivers that are acquired from other hardware vendors can be used if they are associated with the GENERICTAPE device class. Generic device drivers are not supported in write-one read-many (WORM) device classes.

Linux

You can use the IBM Storage Protect Passthru device driver. IBM Storage Protect Passthru device drivers require the Linux SCSI generic (sg) device driver along with the Linux operating system to install the kernels.

For example, you can install the IBM Storage Protect Passthru device driver for the following tape devices:

ECART
LTO (not from IBM)

All SCSI-attached libraries that contain tape drives that are not IBM labeled from the list must also use the IBM Storage Protect Passthru device driver.

You cannot use the generic SCSI tape (st) device driver that is provided by the Linux operating system. Therefore, the GENERICTAPE device type is not supported for the **DEFINE DEVCLASS** command.

Windows You can select a Windows Hardware Qualification Lab certified device driver instead of the IBM Storage Protect device driver. The Windows Hardware Qualification Lab certified device driver can be used only for devices that have a non-IBM label and for non-IBM tape drives. For the Windows Hardware Qualification Lab certified device driver, you can select either the IBM Storage Protect SCSI passthru device driver or the Windows tape device driver. If the SCSI passthru device driver is used, the device class on the **DEFINE DEVCLASS** command cannot be GENERICTAPE. If the Windows tape device driver is used, the device class must be GENERICTAPE.

Special file names for tape devices

A tape device must have a special file name for the server to work with tape, medium changer, or removable media devices.

AIX

When a device is configured successfully, a logical file name is returned. Table 19 on page 78 specifies the name of the device, also called a special file name, that corresponds to the drive or library. You can use the **SMIT** operating system command to get the device special file name. In the examples, x specifies an integer, 0 or greater.

Table 19. Device examples

Device	Device example	Logical file name
Tape drives that can be used by the IBM Storage Protect device driver.	/dev/mtx	mtx
Tape drives that can be used by the IBM tape device driver.	/dev/rmtx	rmtx
Tape drives that can be used by the IBM AIX generic tape device driver.	/dev/rmtx	rmtx
Library devices that can be used by the IBM Storage Protect device driver.	/dev/lbx	lbx
Library devices that can be used by the IBM tape device driver.	/dev/smcx	smcx

Linux

When a device is configured successfully, a logical file name is returned. Table 20 on page 78 specifies the name of the device, also called the special file name, that corresponds to the drive or library. In the examples, x specifies an integer, 0 or greater.

Table 20. Device examples

Device	Device example	Logical file name
Tape drives that can be used by the IBM Storage Protect passthru device driver.	/dev/tsm SCSI /mtx	mtx
Tape drives that can be used by the IBM lin_tape device driver.	/dev/IBMtape x	IBMtape x

Table 20. Device examples (continued)

Device	Device example	Logical file name
Library devices that can be used by the IBM Storage Protect passthru device driver.	/dev/tmscsi /lb x	lb x
Library devices that can be used by the IBM lin_tape device driver.	/dev/IBMchan ger x	IBMchanger x

Windows

When a device is configured successfully, a logical file name is returned. Table 21 on page 79 specifies the name of the device, also called the special file name, that corresponds to the drive or library. In the examples, a , b , c , d , and x specify an integer, 0 or greater, where:

- a specifies the target ID.
- b specifies the LUN.
- c specifies the SCSI bus ID.
- d specifies the port ID.

Table 21. Device examples

Device	Device example	Converted device name
Tape drives that are supported by the IBM Storage Protect device driver.	mta. $b.c.d$	mta. $b.c.d$
Tape drives that are supported by the IBM Storage Protect passthru device driver.	mta. $b.c.d$	mta. $b.c.d$
Tape drives that are supported by the IBM device driver.	Tapex	mta. $b.c.d$
Library devices that are supported by the IBM Storage Protect device driver.	lba. $a.b.c.d$	lba. $b.c.d$
Library devices that are supported by the IBM Storage Protect passthru device driver.	lba. $b.c.d$	lba. $b.c.d$
Library devices that are supported by the IBM device driver.	Changerx	lba. $b.c.d$

Installing and configuring tape device drivers

Before you can use tape devices with IBM Storage Protect, you must install the correct tape device driver.

IBM Storage Protect supports all devices that are supported by IBM tape device drivers. However, IBM Storage Protect does not support all the operating-system levels that are supported by IBM tape device drivers.

Installing and configuring IBM device drivers for IBM tape devices

Install and configure an IBM tape device driver to use an IBM tape device.

About this task

For instructions about installing and configuring IBM tape device drivers, see the [IBM Tape Device Drivers Installation and User's Guide](#).

AIX After you complete the installation procedure in the *IBM Tape Device Drivers Installation and User's Guide*, different messages are issued, depending on the device driver that you are installing. If you are installing the device driver for an IBM tape drive or library, the following messages are returned:

```
rmtx Available
```

or

```
smcx Available
```

Note the value of *x*, which is assigned by the IBM tape device driver. To determine the special file name of your device, issue one of the following commands:

- For tape drives, `ls -l /dev/rmt*`
- For tape libraries, `ls -l /dev/smc*`

The file name might have more characters at the end to indicate different operating characteristics, but these characters are not needed by IBM Storage Protect. For IBM device drivers, use the base file name in the **DEVICE** parameter of the **DEFINE PATH** command to assign a device to a drive (`/dev/rmtx`) or a library (`/dev/smcx`).

After you install the device driver, you can use the System Management Interface Tool (SMIT) to configure non-IBM tape drives and tape libraries. Complete the following steps:

1. Run the SMIT program.
2. Click **Devices**.
3. Click **IBM Storage Protect Devices**.
4. Click **Fibre Channel SAN Attached devices**.
5. Click **Discover Devices Supported by IBM Storage Protect**. Wait for the discovery process to be completed.
6. Go back to the **Fibre Channel SAN Attached devices** menu, and click **List Attributes of a Discovered Device**.

Linux After you complete the installation procedure in the *IBM Tape Device Drivers Installation and User's Guide*, different messages are issued, depending on the device driver that you are installing. If you are installing the device driver for an IBM LTO or 3592 device, the following messages are returned:

```
IBMtapex Available
```

or

```
IBMChangerx Available
```

Note the value of *x*, which is assigned by the IBM tape device driver. To determine the special file name of your device, issue one of the following commands:

- For tape drives, `ls -l /dev/IBMtape*`
- For tape libraries, `ls -l /dev/IBMChanger*`

The file name might have more characters at the end to indicate different operating characteristics, but these characters are not needed by IBM Storage Protect. For IBM device drivers, use the base file name in the **DEVICE** parameter of the **DEFINE PATH** command to assign a device to a drive (`/dev/IBMtapex`) or a library (`/dev/IBMChangerx`).

Restriction: The device type of this class must not be **GENERICTAPE**.

Windows For Windows operating systems, IBM Storage Protect provides two device drivers:

Passthru device driver

If the tape device manufacturer provides a SCSI device driver, install the IBM Storage Protect passthru device driver.

SCSI device driver for tape devices

If the tape device manufacturer does not provide a SCSI device driver, install the IBM Storage Protect SCSI device driver for tape devices. The driver file name is `tsmscsi64.sys`.

For instructions about installing and configuring IBM tape device drivers, see the *IBM Tape Device Drivers Installation and User's Guide*. After you install the IBM tape device driver, the server specifies a special file name, `TapeX`, for IBM tape drives, or `ChangerY`, for IBM medium changers. For an IBM Storage Protect SCSI device driver or an IBM Storage Protect passthru device driver, you can issue the Windows operating system command, **regedit**, to verify the device special file name and driver. The IBM Storage Protect server also provides a utility to check the device for the Windows operating system. The utility, **tsmdlst**, is packaged with the server package. To use the utility, complete the following steps:

1. Ensure that the host bus adapter application programming interface (API) is installed.
2. To obtain device information from the host system, type:

```
tsmdlst
```

Related concepts

Multipath I/O access with IBM tape devices

Multipath I/O is a technique that uses different paths to access the same physical device, for example through multiple host bus adapters (HBA) or switches. The use of the multipath technique helps to ensure that a single point of failure does not occur.

Multipath I/O access with IBM tape devices

Multipath I/O is a technique that uses different paths to access the same physical device, for example through multiple host bus adapters (HBA) or switches. The use of the multipath technique helps to ensure that a single point of failure does not occur.

The IBM tape device driver provides multipathing support so that if one path fails, the server can use a different path to access data on a storage device. The failure and transition to a different path are undetected by the running server or by a storage agent. The IBM tape device driver also uses multipath I/O to provide dynamic load balancing for enhanced I/O performance.

To provide redundant paths for IBM tape devices, connect each device in one of the following configurations:

- Connect to two or more ports on a multiport Fibre Channel.
- Connect to an SAS Host Bus Adapter, if available on your operating system.
- Connect to different single Fibre Channel Host Bus Adapters.

If multipath I/O is enabled and a permanent error occurs on one path, such as a malfunctioning HBA or cable, device drivers provide automatic path failover to an alternative path.

After multipath I/O is enabled, the IBM tape device driver detects all paths for a device on the host system. One path is designated as the primary path. The rest of the paths are alternative paths. The maximum number of alternative paths for a device is 16. For each path, the IBM tape device driver creates a special file with a unique name. A path must exist on the system before the driver can create a special file for the path. If a path does not exist, the driver does not create a special file. When you use the **DEFINE PATH** command to specify the path to a destination, specify the file that is associated with the primary path as the value of the **DEVICE** parameter.

AIX

On AIX, multipath I/O is not enabled automatically when the IBM tape device driver is installed. You must configure it for each logical device after installation. Multipath I/O remains enabled until the device is deleted or the support is unconfigured. For configuration instructions, see the *IBM Tape Device Drivers Installation and User's Guide*.

To obtain the names of special files, use the **ls -l** command, for example, **ls -l /dev/rmt***. Primary paths and alternative paths are identified by **PRI** and **ALT**, as seen in the following example:

rmt0 Available 20-60-01-PRI IBM 3590 Tape Drive and Medium Changer (FCP)
rmt1 Available 30-68-01-ALT IBM 3590 Tape Drive and Medium Changer (FCP)

In this example, the following paths are associated with the IBM 3590 tape drive:

- 20-60-01-PRI
- 30-68-01-ALT

The name of the special file that is associated with the primary path is `/dev/rmt0`. Specify `/dev/rmt0` as the value of the **DEVICE** parameter in the **DEFINE PATH** command.

To display path-related details about a particular tape drive, you can also use the **itdt -f /dev/rmtx path** command, where x is the number of the configured tape drive. To display path-related details about a particular medium changer, use the **itdt -f /dev/smcy path** command, where y is the number of the configured medium changer.

Linux

On Linux, multipath I/O for medium changers and tape drives is not enabled automatically when the device driver is installed. For instructions about configuring multipath I/O, see the *IBM Tape Device Drivers Installation and User's Guide*.

When multipath I/O is enabled for a logical device, it remains enabled until the device is deleted or the support is unconfigured.

To display the names of special file for IBM tape drives and medium changers, use the **ls -l /dev/IBMx** command, where x is the index number of the device. You can also enter the **cat /proc/scsi/IBMtape** command for tape drives. As shown in the `IBMTape` file, primary paths, and alternative paths are identified as Primary or Alternate:

Number	Model	SN	HBA	FO Path
0	03592	IBM1234567	qla2xxx	Primary
1	03592	IBM1234567	qla2xxx	Alternate

The name of the special file that is associated with the primary path for this tape drive is `/dev/IBMTape0`. Specify `/dev/IBMTape0` as the value of the **DEVICE** parameter in the **DEFINE PATH** command for this device.

To obtain the names of the special files that are associated with the primary paths for all medium changers that are configured on the system, run the **cat /proc/scsi/IBMchanger** command. The following example is taken from the `IBMchanger` file:

Number	Model	SN	HBA	FO Path
3	03584L22	IBM1002345	qla2xxx	Primary
4	03584L22	IBM1002345	qla2xxx	Alternate

The name of the special file that is associated with the primary path for this medium changer is `/dev/IBMchanger3`. Specify `/dev/IBMchanger3` as the value of the **DEVICE** parameter in the **DEFINE PATH** command for this device.

To display path-related details about a particular tape drive on the system, use the **itdt -f /dev/IBMtapex path** command, where x is the number of a configured tape device. To display path-related details about a particular medium changer on the system, use the **itdt -f /dev/IBMchangerx path** command, where x is the number of a configured medium changer.

Windows

On Windows, multipath I/O for medium changers and tape drives is not enabled automatically when the device driver is installed. For instructions about configuring multipath I/O, see the *IBM Tape Device Drivers Installation and User's Guide*. If multipath I/O is configured, a device has two matching device names with different locations. To obtain detailed information about the primary path and the alternative path, run the IBM Tape Diagnostic Tool with the **qrypath** function. The output is similar to the following example:

```
C:\Users\Administrator\Downloads\ITDT> .\itdt.exe qrypath -f \\.\Tape0
Querying SCSI paths...
```

```

Total paths configured..... 2

Alternate Path
Logical Device..... Tape0
Serial Number..... 0000078F7612
SCSI Host ID..... 8
SCSI Channel..... 0
Target ID..... 3
Logical Unit..... 0
Path Enabled..... Yes

Primary Path
Logical Device..... Tape0
Serial Number..... 0000078F7612
SCSI Host ID..... 8
SCSI Channel..... 0
Target ID..... 1
Logical Unit..... 0
Path Enabled..... Yes

Exit with code: 0

```

AIX Configuring tape device drivers on AIX systems

Review the instructions to install and configure non-IBM tape device drivers on AIX systems.

About this task

For instructions about installing and configuring IBM tape device drivers, see the [IBM Tape Device Drivers Installation and User's Guide](#).

AIX SCSI and Fibre Channel devices

The IBM Storage Protect device definition menus and prompts in SMIT allow for the management of both SCSI and Fibre Channel (FC) attached devices.

The main menu for IBM Storage Protect has two options:

SCSI attached devices

Use this option to configure SCSI devices that are connected to a SCSI adapter in the host.

Fibre channel system area network (SAN) attached devices

Use this option to configure devices that are connected to an FC adapter in the host. Choose one of the following attributes:

List attributes of a discovered device

Lists attributes of a device that is known to the current ODM database.

- FC Port ID:

The 24-bit FC Port ID(N(L)_Port or F(L)_Port). This is the address identifier that is unique within the associated topology where the device is connected. In the switch or fabric environments, it can be determined by the switch, with the upper 2 bytes, which are not zero. In a Private Arbitrated Loop, it is the Arbitrated Loop Physical Address(AL_PA), with the upper 2 bytes being zero. Consult with your FC vendors to find out how an AL_PA or a Port ID is assigned.

- Mapped LUN ID:

An FC to SCSI bridge (also, called a converter, router, or gateway) box. Consult with your bridge vendors about how LUNs are mapped. You should not change LUN Mapped IDs.

- WW Name:

The worldwide name of the port to which the device is attached. It is the 64-bit unique identifier that is assigned by vendors of FC components such as bridges or native FC devices. Consult with your FC vendors to find out the WWN of a port.

- Product ID:

The product ID of a device. Consult with your device vendors to determine the product ID.

Discover devices supported by IBM Storage Protect

This option discovers devices on an FC SAN that are supported by IBM Storage Protect and makes them available. If a device is added to or removed from an existing SAN environment, rediscover devices by selecting this option. Devices must be discovered first so that current values of device attributes are shown in the **List Attributes of a Discovered Device** option. Supported devices on FC SAN are tape drives, and autochangers. The IBM Storage Protect device driver ignores all other device types, such as disk.

Remove all defined devices

This option removes all FC SAN-attached IBM Storage Protect devices whose state is **DEFINED** in the ODM database. If necessary, rediscover devices by selecting the **Discover Devices Supported by IBM Storage Protect** option after the removal of all defined devices.

Remove a device

This option removes a single FC SAN-attached IBM Storage Protect device whose state is **DEFINED** in the ODM database. If necessary, rediscover the device by selecting the **Discover Devices Supported by IBM Storage Protect** option after removal of a defined device.

AIX Configuring IBM Storage Protect device drivers for autochangers

Use the following procedure to configure IBM Storage Protect device drivers for autochangers for non-IBM libraries.

Procedure

Run the SMIT program to configure the device driver for each autochanger or robot:

1. Select **Devices**.
2. Select **IBM Storage Protect Devices**.
3. Select **Library/MediumChanger**.
4. Select **Add a Library/MediumChanger**.
5. Select the IBM Storage Protect-SCSI-LB for any IBM Storage Protect supported library.
6. Select the parent adapter to which you are connecting the device. This number is listed in the form: 00-0X, where X is the slot number location of the SCSI adapter card.
7. When prompted, enter the **CONNECTION** address of the device that you are installing. The connection address is a two-digit number. The first digit is the SCSI ID (the value you recorded on the worksheet). The second digit is the device's SCSI logical unit number (LUN), which is usually zero, unless otherwise noted. The SCSI ID and LUN must be separated by a comma (,).
For example, a connection address of 4,0 has a SCSI ID=4 and a LUN=0.
8. Click **DO**.

You receive a message (logical file name) of the form **lbX Available**. Note the value of X, which is a number that is assigned automatically by the system. Use this information to complete the **Device Name** field on your worksheet.

For example, if the message is **lb0 Available**, the **Device Name** field is **/dev/lb0** on the worksheet. Always use the **/dev/** prefix with the name provided by SMIT.

AIX Configuring IBM Storage Protect device drivers for tape drives

Use the following procedure to configure IBM Storage Protect device drivers for autochangers for vendor-acquired libraries.

Procedure

Important: IBM Storage Protect cannot overwrite *tar* or *dd* tapes, but *tar* or *dd* can overwrite IBM Storage Protect tapes.

Restriction: Tape drives can be shared only when the drive is not defined or the server is not started. The **MKSYSB** command does not work when both IBM Storage Protect and AIX are sharing the same drive

or drives. To use the operating system's native tape device driver with a SCSI drive, the device must be configured to AIX first and then configured to IBM Storage Protect. See your AIX documentation regarding these native device drivers.

Run the SMIT program to configure the device driver for each drive (including drives in libraries) as follows:

1. Select **Devices**.
2. Select **IBM Storage Protect Devices**.
3. Select **Tape Drive**.
4. Select **Add a Tape Drive**.
5. Select the IBM Storage Protect-SCSI-MT for any supported tape drive.
6. Select the adapter to which you are connecting the device. This number is listed in the form: 00-0X, where X is the slot number location of the SCSI adapter card.
7. When prompted, enter the CONNECTION address of the device you are installing. The connection address is a two-digit number. The first digit is the SCSI ID (the value you recorded on the worksheet). The second digit is the device's SCSI logical unit number (LUN), which is usually zero, unless otherwise noted. The SCSI ID and LUN must be separated by a comma (,).
For example, a connection address of 4,0 has a SCSI ID=4 and a LUN=0.
8. Click **DO**. You receive a message:

If you are configuring the device driver for a tape device (other than an IBM tape drive), you receive a message (logical file name) of the form `mtX Available`. Note the value of X, which is a number that is assigned automatically by the system. Use this information to complete the **Device Name** field on the worksheet.

For example, if the message is `mt0 Available`, the **Device Name** field is `/dev/mt0` on the worksheet. Always use the `/dev/` prefix with the name provided by SMIT.

AIX Configuring Fibre Channel SAN-attached devices

To configure a Fibre Channel SAN-attached device, complete the procedure.

Procedure

1. Run the SMIT program.
2. Select **Devices**.
3. Select **IBM Storage Protect Devices**.
4. Select **Fibre Channel SAN Attached devices**.
5. Select **Discover Devices Supported by IBM Storage Protect**. The discovery process can take some time.
6. Go back to the **Fibre Channel** menu, and select **List Attributes of a Discovered Device**.
7. Note the three-character device identifier, which you use when you define a path to the device to IBM Storage Protect.
For example, if a tape drive has the identifier `mt2`, specify `/dev/mt2` as the device name.

Linux Configuring tape device drivers on Linux systems

Review the following topics when you install and configure tape device drivers on Linux systems.

Linux **Configuring IBM Storage Protect passthru drivers for tape devices and libraries**

To use the IBM Storage Protect Linux Passthru driver, you must complete the following steps.

Procedure

1. Verify that the device is connected to your system, and is powered on and active.
2. Verify that the device is correctly detected by your system by issuing this command:

```
cat /proc/scsi/scsi
```

3. Ensure that both the IBM Storage Protect device driver package (tsmscsi) and the storage server package are installed.
4. There are two driver configuration methods available in the IBM Storage Protect device driver package: `autoconf` and `tsmscsi`. Both of these methods complete the following tasks:
 - Load the Linux SCSI generic driver (`sg`) to the kernel.
 - Create necessary special files for the Passthru driver.
 - Create device information files for tape devices (`/dev/tsmscsi/mtinfo`) and libraries (`/dev/tsmscsi/lbinfo`).
5. Run the configuration method that you prefer (`autoconf` or `tsmscsi`) for the IBM Storage Protect Passthru driver.
 - To run the `autoconf` configuration method, issue the following command:

```
autoconf
```

- To run the `tsmscsi` configuration method, complete the following steps:
 - a. Copy the two sample configuration files that are in the installation directory from `mt.conf.smp` and `lb.conf.smp` to `mt.conf` and `lb.conf`, respectively.
 - b. Edit the `mt.conf` and `lb.conf` files. Add one stanza (as shown in the example at the start of the file) for each SCSI target, ID, and LUN combination. Each combination of SCSI target, ID, and LUN entries correspond to a tape drive or library you want configured. Make sure that the files meet these requirements:
 - Remove the example that is at the start of the files.
 - There must be a new line between each stanza.
 - There must be one new line after the last stanza.
 - Ensure that there are no number signs (#) in either file.
 - c. Run the `tsmscsi` script from the device driver installation directory.
6. Verify that the device is configured properly by viewing the text files for tape devices (`/dev/tsmscsi/mtinfo`) and libraries (`/dev/tsmscsi/lbinfo`).
 7. Determine the special file names for the tape drives and libraries:
 - To determine the names for tape devices, issue the following command:

```
> ls /dev/tsmscsi/mt*
```

- To determine the names for libraries, issue the following command:

```
> ls /dev/tsmscsi/lb*
```

This information helps you identify which of the `/dev/tsmscsi/mtx` and `/dev/tsmscsi/lbx` special file names to provide the server when you issue a **DEFINE PATH** command.

What to do next

If you restart the host system, you must rerun the `autoconf` or `tsmscsi` script to reconfigure IBM Storage Protect devices. If you restart the IBM Storage Protect server instance, you do not have to reconfigure devices. In general, the Linux SCSI generic driver is preinstalled to the kernel. To verify that the driver is in the kernel, issue the following command:

```
> lsmod | grep sg
```

If the driver is not in the kernel, issue the `modprobe sg` command to load the `sg` driver into the kernel.

Linux Installing zSeries Linux Fibre Channel adapter (zfcp) device drivers

The zSeries Linux Fibre Channel adapter (zfcp) device driver is a special adapter driver on the IBM zSeries system.

About this task

IBM Storage Protect and IBM tape device drivers can run on zSeries platforms with Linux operating systems in 64-bit environments, and support most original equipment manufacturer (OEM) and IBM tape devices with Fibre Channel interfaces.

For more information about the `zfcp` driver, see the IBM Redpaper, *Getting Started with zSeries Fibre Channel Protocol*, which is available at [IBM Redbooks®](#).

Procedure

1. Load the `qdio` module.
2. Install the `zfcp` driver.
3. Map the Fibre Channel Protocol (FCP) and configure the `zfcp` driver.
4. Install and configure the IBM tape device driver.

Linux Information about your system's SCSI devices

Information about the devices seen by your system is available in the file `/proc/scsi/scsi`. This file contains a list of every detected SCSI device.

The following device information is available: the host number, channel number, SCSI ID, Logical Unit number, vendor, firmware level, type of device, and the SCSI mode. For example, if a system contains some StorageTek and IBM libraries, a SAN Gateway, and some Quantum DLT drives, the `/proc/scsi/scsi` file will look similar to this:

```
Attached devices:
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: STK      Model: 9738      Rev: 2003
  Type:  Medium Changer          ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: PATHLIGHT Model: SAN Gateway Rev: 32aC
  Type: Unknown                 ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: QUANTUM  Model: DLT7000    Rev: 2560
  Type: Sequential-Access       ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 01 Lun: 04
  Vendor: IBM      Model: 7337      Rev: 1.63
  Type: Medium Changer          ANSI SCSI revision: 02
```

Linux Preventing tape labels from being overwritten

The IBM Storage Protect Passthru device driver uses the Linux SCSI generic device driver (`sg`) to control and operate tape devices that are attached on the system. If the Linux generic SCSI tape device driver (`st`) is loaded to the kernel and configures attached tape devices, conflicts can arise over how a device is managed because the generic `sg` driver and the `st` driver can both control the same device.

About this task

If the st driver controls devices that are used by IBM Storage Protect, IBM Storage Protect internal tape labels can be overwritten and data can be lost. If an application uses the st driver to control devices and the non-rewind option is not specified, tapes are automatically rewound following completion of an operation. The auto-rewind operation relocates the tape header position to the beginning of the tape. If the tape remains loaded in the drive, the next non-IBM Storage Protect write operation overwrites the IBM Storage Protect tape label because the label is at the beginning of the tape.

To prevent IBM Storage Protect labels from being overwritten, which can result in data loss, ensure that only the IBM Storage Protect Passthru driver controls devices that are used by IBM Storage Protect. Remove the st driver from the kernel or, if the driver is used by some applications on the system, delete the special files that correspond to IBM Storage Protect devices so that the st driver can no longer control them.

If you are using the IBM tape device driver to control devices on your system, you might encounter the same issues with device driver control conflicts. Review your IBM tape documentation to determine how to resolve this issue and prevent data loss.

Remove the st driver

If no other applications on the system use st devices, remove the st driver from the kernel. Issue the following command to unload the st driver:

```
rmmod st
```

Delete device special files that correspond to IBM Storage Protect devices

If there are applications that require use of the st driver, delete the special files that correspond to IBM Storage Protect devices. These special files are generated by the st driver. When they are eliminated, the st driver can no longer control the corresponding IBM Storage Protect devices. Device special file names for tape drives appear in the /dev/ directory. Their names have the form /dev/[n]st[0-1024][1][m][a].

List the st drive special file names and IBM Storage Protect device special file names by using the ls command. Based on the output of the device sequences, you can find devices in the st devices list matching those in the IBM Storage Protect devices list. The rm command can then be used to delete st devices.

Issue the following commands to list the st and IBM Storage Protect devices:

```
ls -l /dev/*st*  
ls -l /dev/tsm SCSI/mt*
```

Delete the st devices with the rm command:

```
rm /dev/*st*
```

Windows Configuring tape device drivers on Windows systems

Review the instructions to install and configure drivers for tape devices and libraries on Windows systems.

Windows Preparing to use the IBM Storage Protect passthru driver for tape devices and libraries

To use the IBM Storage Protect Windows passthru device driver for tape devices and libraries, you must install the driver and obtain the device names for the server to use.

Before you begin

1. Determine whether the manufacturer of the tape device or tape library provides a device driver.
2. If the manufacturer provides a device driver package, download the package and install it.
3. Configure the SCSI device driver by following the manufacturer's instructions.

Procedure

1. Install the IBM Storage Protect passthru device driver.
2. Obtain the device names that the server must use by taking one of the following actions:
 - On the server, run the **QUERY SAN** command. The output shows all devices names and their associated device serial numbers.
 - In the server directory, run the **tsmdlst.exe** utility. The output shows all devices names, their associated serial numbers, and associated device locations.
 - At the Windows system command prompt, run the **regedit** command. From the output, obtain the device file names based on the device locations. The location consists of the port ID, SCSI bus ID, LUN ID, and SCSI target ID. The IBM Storage Protect device file name has a format of **mtA.B.C.D** for tape drives and **lbA.B.C.D** for tape libraries, where:
 - *A* is the SCSI target ID.
 - *B* is the LUN ID.
 - *C* is the SCSI bus ID.
 - *D* is the port ID.

Configuring the IBM Storage Protect SCSI driver for tape devices and libraries

If the manufacturer of a tape drive or tape library does not provide a SCSI device driver, you must install the IBM Storage Protect SCSI device driver.

About this task

The IBM Storage Protect SCSI device driver file name is **tsmscsi64.sys**.

Procedure

1. Locate the device in the Device Manager console (**devmgmt.msc**) and select it. Tape drives are listed under **Tape drives**, and medium changers are under **Medium Changers**.
2. Configure the device for use by the **tsmscsi64.sys** device driver:
 - a. Right-click the device and click **Update Driver Software**.
 - b. Click **Browse my computer for driver software**.
3. Click **Let me pick from a list of device drivers on my computer**.
4. Click **Next**.
5. Select the appropriate option:
 - a. For a tape drive, select **IBM Storage Protect for Tape Drives**.
 - b. For a medium changer, select **IBM Storage Protect for Medium Changers**.
6. Click **Next**.
7. Click **Close**.
8. Verify that the device was configured correctly for the **tsmscsi64** device driver:
 - a. Right-click on the device and click **Properties**.
 - b. Click the **Driver** tab and **Driver Details**. The **Driver Details** window shows the device driver that is controlling the device.

Configuring libraries for use by a server

To use a library or libraries for storage for an IBM Storage Protect server, you must first set up the devices on the server system.

Before you begin

1. Attach devices to the server hardware. Follow the instructions in [“Attaching an automated library device to your system” on page 76](#).
2. Select the tape device drivers. Follow the instructions in [“Selecting a tape device driver” on page 76](#).
3. Install and configure the tape device drivers. Follow the instructions in [“Installing and configuring tape device drivers” on page 79](#).
4. Determine the device names that are needed to define the library to the server. Follow the instructions in [“Special file names for tape devices” on page 78](#).

Procedure

1. Define the library and the path from the server to the library. Follow the instructions in [“Defining libraries” on page 91](#).
2. Define the drives in the library. Follow the instructions in [“Defining drives” on page 92](#).

For SCSI libraries, you can use the **PERFORM LIBACTION** command to define drives and paths for a library in one step, instead of completing both steps [“2” on page 90](#) and [“3” on page 90](#). To use the **PERFORM LIBACTION** command to define drives and paths for a library, the SANDISCOVERY option must be supported and enabled.

3. Define a path from the server to each drive by using the **DEFINE PATH** command.
4. Define a device class. Follow the instructions in [“Defining tape device classes” on page 93](#).

Device classes specify the recording formats for drives and classify them according to type. Use the default value, **FORMAT=DRIVE** as the recording format only if all the drives that are associated with the device class can read and write to all of the media.

For example, you have a mix of Ultrium Generation 3 and Ultrium Generation 4 drives, but you have only Ultrium Generation 3 media. You can specify **FORMAT=DRIVE** because both the Generation 4 and Generation 3 drives can read from and write to Generation 3 media.

5. Define a storage pool by using the **DEFINE STGPOOL** command.

Consider the following key choices for defining storage pools:

- Scratch volumes are empty volumes that are available for use. If you specify a value for the maximum number of scratch volumes in the storage pool, the server can choose from the scratch volumes available in the library.

If you do not allow scratch volumes, you must complete the extra step of explicitly defining each volume to be used in the storage pool. Also, specify the **MAXSCRATCH=0** parameter when you define the storage pool so that scratch volumes are not used.

- The default setting for primary storage pools is collocation by group. The default for copy storage pools and active-data pools is disablement of collocation. The server uses *collocation* to keep all files that belong to a group of client nodes, a single client node, a client file space, or a group of client file spaces on a minimal number of volumes. If collocation is disabled for a storage pool and clients begin storing data, you cannot easily change the data in the pool so that it is collocated.
6. Check in and label library volumes. Follow the instructions in [“Checking volumes into an automated library” on page 176](#) and [“Labeling tape volumes” on page 174](#).

Ensure that enough volumes in the library are available to the server. Keep enough labeled volumes on hand so that you do not run out during an operation such as client backup. Label extra scratch volumes for any potential recovery operations that you might have later.

The procedures for checking in and labeling volumes are the same whether the library contains drives of a single device type, or drives of multiple device types. You can use the **CHECKIN LIBVOLUME** command to check in volumes that are already labeled. Or, if you want to label and check in volumes with one step, issue the **LABEL LIBVOLUME** command.

Libraries with multiple device types: If your library has drives of multiple device types, and you defined two libraries to the IBM Storage Protect server, the two defined libraries represent one physical library. You must check in tape volumes separately to each defined library. Ensure that you check in volumes to the correct IBM Storage Protect library.

What to do next

Verify your device definitions to ensure that everything is configured correctly. Use a **QUERY** command to review information about each storage object.

When you review the results of the **QUERY DRIVE** command, verify that the device type for the drive is what you expect. If a path is not defined, the drive device type is listed as UNKNOWN and if the wrong path is used, GENERIC_TAPE or another device type is shown. This step is especially important when you are using mixed media.

Optionally, configure library sharing. Follow the instructions in [“Configuring library sharing” on page 100](#).

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

[PERFORM LIBACTION \(Define or delete all drives and paths for a library\)](#)

Defining tape devices

Before you can back up or migrate data to tape, you must define a tape device to the server.

Defining libraries and drives

A tape library can include one or more tape drives. Learn how to define libraries, drives, and paths to the IBM Storage Protect server.

Defining libraries

Before you can use a drive, you must define the library to which the drive belongs.

Procedure

1. Define the library by using the **DEFINE LIBRARY** command.

For example, if you have an IBM TS3500 tape library, you can define a library that is named ROBOTMOUNT by using the following command:

```
define library robotmount libtype=scsi
```

If you require library sharing or LAN-free data movement, see the following information:

- [“Configuring library sharing” on page 100](#)
 - [“Configuring LAN-free data movement” on page 118](#)
2. Define a path from the server to the library by using the **DEFINE PATH** command. When you specify the **DEVICE** parameter, enter the device special file name. This name is required by the server to

communicate with tape drives, medium changer, and removable media devices. For more information about device special file names, see [“Special file names for tape devices”](#) on page 78.

```
AIX define path server1 robotmount srctype=server desttype=library
device=/dev/lb0
```

```
Linux define path server1 robotmount srctype=server desttype=library
device=/dev/tsm SCSI/lb0
```

```
Windows define path server1 robotmount srctype=server desttype=library
device=lb0.0.1.0
```

Related information

[DEFINE LIBRARY](#) (Define a library)

[DEFINE PATH](#) (Define a path)

Defining SCSI libraries on a SAN

For a library type of SCSI on a SAN, the server can track the library's serial number. With the serial number, the server can confirm the identity of the device when you define the path or when the server uses the device.

About this task

If you choose, you can specify the serial number when you define the library to the server. For convenience, the default is to allow the server to obtain the serial number from the library when you define the path.

If you specify the serial number, the server confirms that the serial number is correct when you define the path to the library. When you define the path, you can set the **AUTODETECT=YES** parameter to allow the server to correct the serial number if the number that it detects does not match what you entered when you defined the library. As a best practice, specify the **AUTODETECT=YES** parameter to automatically update the serial number for the drive in the database when the path is defined.

Depending on the capabilities of the library, the server might not be able to automatically detect the serial number. Not all devices are able to return a serial number when prompted by an application such as the server. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device. For more information, see [“Impacts of device changes on the SAN”](#) on page 128.

Defining drives

To inform the server about a drive that can be used to access storage volumes, issue the **DEFINE DRIVE** command, followed by the **DEFINE PATH** command.

Before you begin

A *drive object* represents a drive mechanism within a library that uses removable media. For devices with multiple drives, including automated libraries, you must define each drive separately and associate it with a library. Drive definitions can include such information as the element address for drives in SCSI, how often a tape drive is cleaned, and whether the drive is online.

IBM Storage Protect supports tape drives that can be stand-alone or that can be part of an automated library. The preferred method is to configure the tape solution by using automated libraries.

About this task

When you issue the **DEFINE DRIVE** command, you must provide some or all of the following information:

Library name

The name of the library in which the drive is located.

Drive name

The name that is assigned to the drive.

Serial number

The serial number of the drive. The serial number parameter applies only to drives in SCSI. With the serial number, the server can confirm the identity of the device when you define the path or when the server uses the device.

You can specify the serial number if you choose. The default is to enable the server to obtain the serial number from the drive itself at the time that the path is defined. If you specify the serial number, the server confirms that the serial number is correct when you define the path to the drive. When you define the path, you can set the **AUTODETECT=YES** parameter to enable the server to correct the serial number if the number that it detects does not match what you entered when you defined the drive. As a best practice, specify the **AUTODETECT=YES** parameter to automatically update the serial number for the drive in the database when the path is defined.

Depending on the capabilities of the drive, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device. See [“Impacts of device changes on the SAN” on page 128](#).

Element address

The element address of the drive. The **ELEMENT** parameter applies only to drives in SCSI libraries. The element address is a number that indicates the physical location of a drive within an automated library. The server needs the element address to connect the physical location of the drive to the drive's SCSI address. The server can obtain the element address from the drive when you define the path, or you can specify the element number when you define the drive. As a best practice, specify the **ELEMENT=AUTODETECT** parameter for the server to automatically detect the element number when the path to the drive is defined.

Depending on the capabilities of the library, the server might not be able to automatically detect the element address. In this case, you must supply the element address when you define the drive, if the library has more than one drive. To obtain the element address, go to the [IBM Support Portal for IBM Storage Protect](#).

Tip: IBM tape device drivers and non-IBM tape device drivers generate different device files and formats:

- For IBM, device names begin with `mt` followed by an integer, for example, `/dev/mt0`.
- For IBM Storage Protect tape device drivers, tape device names begin with `mt` followed by an integer, for example `/dev/mt0`.

You must use the correct device file when you define a path.

Procedure

1. Assign a drive to a library by issuing the **DEFINE DRIVE** command.
2. To make the drive usable by the server, issue the **DEFINE PATH** command.

For examples about configuring libraries, paths, and drives, see [Example: Configure a SCSI or virtual tape library with a single drive device type](#) and [Example: Configure a SCSI or virtual tape library with multiple drive device types](#).

Defining tape device classes

A device class defines a set of characteristics that are used by a set of volumes that can be created in a storage pool. You must define a device class for a tape device to ensure that the server can use the device.

Before you begin

You must define libraries and drives to the server before you define device classes.

About this task

For a list of supported devices and valid device class formats, see the IBM Storage Protect Supported Devices website for your operating system:

- [AIX](#) | [Windows](#) Supported devices for AIX and Windows
- [Linux](#) Supported devices for Linux

You can define multiple device classes for each device type. For example, you might want to specify different attributes for different storage pools that use the same type of tape drive. Variations might be required that are not specific to the device, but rather to how you want to use the device (for example, mount retention or mount limit).

Guidelines:

- One device class can be associated with multiple storage pools, but each storage pool is associated with only one device class.
- SCSI libraries can include tape drives of more than one device type. When you define the device class in this environment, you must declare a value for the **FORMAT** parameter.

For more information, see *Mixed device types in libraries* in IBM Documentation.

Procedure

To define a device class, use the **DEFINE DEVCLASS** command with the **DEVTYPE** parameter, which assigns a device type to the device class.

Results

If you include the DEVCONFIG option in the dsmserve.opt file, the files that you specify with that option are automatically updated with the results of the **DEFINE DEVCLASS**, **UPDATE DEVCLASS**, and **DELETE DEVCLASS** commands.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[QUERY DEVCLASS \(Display information on one or more device classes\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Defining LTO device classes

To prevent problems when you mix different generations of LTO drives and media in a single library, review the restrictions. Also, review the restrictions for LTO drive encryption.

Mixing LTO drives and media in a library

When you mix different generations of LTO drives and media, you must consider the read/write capabilities of each generation. The preferred method is to configure a different device class for each generation of media.

About this task

If you are considering mixing different generations of LTO media and drives, review the following restrictions:

Table 22. Read/write capabilities for different generations of LTO drives									
Drives	Generation 1 media	Generation 2 media	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media	Generation 7 media	Generation M8 media	Generation 8 media
Generation 1	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Generation 2	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a	n/a

Table 22. Read/write capabilities for different generations of LTO drives (continued)									
Drives	Generation 1 media	Generation 2 media	Generation 3 media	Generation 4 media	Generation 5 media	Generation 6 media	Generation 7 media	Generation M8 media	Generation 8 media
Generation 3	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a	n/a
Generation 4	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a	n/a
Generation 5	n/a	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a	n/a
Generation 6	n/a	n/a	n/a	Read-only access	Read/write access	Read/write access	n/a	n/a	n/a
Generation 7	n/a	n/a	n/a	n/a	Read access	Read/write access	Read/write access	n/a	n/a
Generation 8	n/a	n/a	n/a	n/a	n/a	n/a	Read/write access	Read/write access	Read/write access

Example

If you are mixing different types of drives and media, configure different device classes: one for each type of media. To specify the media type, use the **FORMAT** parameter in each of the device class definitions. (Do not specify FORMAT=DRIVE.) For example, if you are mixing Ultrium Generation 5 and Ultrium Generation 6 drives, specify FORMAT=ULTRIUM5C (or ULTRIUM5) for the Ultrium Generation 5 device class, and FORMAT=ULTRIUM6C (or ULTRIUM6) for the Ultrium Generation 6 device class.

In this example, both device classes can point to the same library with Ultrium Generation 5 and Ultrium Generation 6 drives. The drives are shared between the two storage pools. One storage pool uses the first device class and Ultrium Generation 5 media exclusively. The other storage pool uses the second device class and Ultrium Generation 6 media exclusively. Because the two storage pools share a single library, Ultrium Generation 5 media can be mounted on Ultrium Generation 6 drives as they become available during mount point processing.

If you mix older read-only media generations with newer read/write media in a single library, you must mark the read-only media as read-only and check out all read-only scratch media. For example, if you are mixing Ultrium Generation 4 with Ultrium Generation 6 drives and media in a single library, you must mark the Generation 4 media as read-only. In addition, you must check out all Generation 4 scratch volumes.

Mount limits in LTO mixed-media environments

In a mixed-media library, in which multiple device classes point to the same library, compatible drives are shared between storage pools. Ensure that you set an appropriate value for the **MOUNTLIMIT** parameter in each of the device classes.

For example, in a mixed media library that contains Ultrium Generation 1 and Ultrium Generation 2 drives and media, Ultrium Generation 1 media can be mounted in Ultrium Generation 2 drives.

Consider the example of a mixed library that consists of the following drives and media:

- Four LTO Ultrium Generation 1 drives and LTO Ultrium Generation 1 media
- Four LTO Ultrium Generation 2 drives and LTO Ultrium Generation 2 media

You created the following device classes:

- LTO Ultrium Generation 1 device class LTO1CLASS specifying FORMAT=ULTRIUM1C
- LTO Ultrium Generation 2 device class LTO2CLASS specifying FORMAT=ULTRIUM2C

You also created the following storage pools:

- LTO Ultrium Generation 1 storage pool LTO1POOL based on device class LTO1CLASS
- LTO Ultrium Generation 2 storage pool LTO2POOL based on device class LTO2CLASS

The number of mount points available for use by each storage pool is specified in the device class by using the **MOUNTLIMIT** parameter. The **MOUNTLIMIT** parameter in the LTO2CLASS device class must be set to 4 to match the number of available drives that can mount only LTO7 media. The **MOUNTLIMIT** parameter in the LTO1CLASS device class must be set to a value that is greater than the number of available drives (5 or possibly 6) to adjust for the fact that Ultrium Generation 1 media can be mounted in Ultrium Generation 7 drives. The optimal value for **MOUNTLIMIT** depends on workload and storage pool access patterns.

Monitor and adjust the **MOUNTLIMIT** setting to suit changing workloads. If the **MOUNTLIMIT** for LTO1POOL is set too high, mount requests for the LTO2POOL might be delayed or fail because the Ultrium Generation 2 drives are used to satisfy Ultrium Generation 1 mount requests. In the worst scenario, too much competition for Ultrium Generation 2 drives might cause mounts for Generation 2 media to fail with the following message:

```
ANR8447E No drives are currently available in the library.
```

If the **MOUNTLIMIT** value for LTO1POOL is not set high enough, mount requests that might be satisfied by LTO Ultrium Generation 2 drives are delayed.

Restriction: Restrictions apply when you mix Ultrium Generation 1 with Ultrium Generation 2 or Generation 3 drives because of how mount points are allocated. For example, processes that require multiple mount points that include both Ultrium Generation 1 and Ultrium Generation 2 volumes might try to reserve Ultrium Generation 2 drives only, even when one mount can be satisfied by an available Ultrium Generation 6 drive. Processes that behave in this manner include the **MOVE DATA** and **BACKUP STGPOOL** commands. These processes wait until the required number of mount points can be satisfied with Ultrium Generation 2 drives.

Related information

[BACKUP STGPOOL \(Back up primary storage pool data to a copy storage pool\)](#)

[DEFINE DEVCLASS \(Define a device class\)](#)

[MOVE DATA \(Move files on a storage pool volume\)](#)

Enabling and disabling drive encryption for LTO Generation 4 or later tape drives

IBM Storage Protect supports the three types of drive encryption that are available with LTO Generation 4 or later drives: Application, System, and Library. These methods are defined through the hardware.

About this task

The **DRIVEENCRYPTION** parameter on the **DEFINE DEVCLASS** command specifies whether drive encryption is allowed for IBM and HP LTO Generation 4 or later, Ultrium 4, and Ultrium 4C formats. This parameter ensures IBM Storage Protect compatibility with hardware encryption settings for empty volumes. You cannot use this parameter for storage pool volumes that are full or are filling.

IBM Storage Protect supports the Application method of encryption with IBM and HP LTO-4 or later drives. Only IBM LTO-4 or later supports the System and Library methods. The Library method of encryption can be used only if your system hardware (for example, IBM TS3500) supports it.

Restriction: You cannot use drive encryption with write-once, read-many (WORM) media.

The Application method is defined through the hardware. To use the Application method, in which IBM Storage Protect generates and manages encryption keys, set the **DRIVEENCRYPTION** parameter to ON. This action enables data encryption for empty volumes. If the parameter is set to ON and the hardware is configured for another encryption method, backup operations fail.

Procedure

The following simplified example shows the steps that you would take to enable and disable data encryption for empty volumes in a storage pool:

1. Define a library by issuing the **DEFINE LIBRARY** command:

```
define library 3584 libtype=SCSI
```

2. Define a device class, LTO_ENCRYPT, by issuing the **DEFINE DEVCLASS** command and specifying IBM Storage Protect as the key manager:

```
define devclass lto_encrypt library=3584 devtype=lto driveencryption=on
```

3. Define a storage pool by issuing the **DEFINE STGPPOOL** command:

```
define stgpool lto_encrypt_pool lto_encrypt
```

4. To disable encryption on new volumes, set the **DRIVEENCRYPTION** parameter to OFF. The default value is ALLOW. Drive encryption for empty volumes is allowed if another method of encryption is enabled.

Related concepts

[Tape encryption methods](#)

Deciding on the encryption method to use depends on how you want to manage your data.

Defining 3592 device classes

Device class definitions for 3592, TS1130, TS1140, TS1150, and later devices include parameters for faster volume-access speeds and drive encryption. To prevent problems when mixing different generations of 3592 and TS1130 and later drives in a library, review the guidelines.

Mixing generations of 3592 drives and media in a single library

For optimal performance, do not mix generations of 3592 media in a single library. Media problems can result when different drive generations are mixed. For example, IBM Storage Protect might not be able to read a volume's label.

About this task

The following table shows read/write interoperability for drive generations.

Drives	Generati on 1 format	Generati on 2 format	Generati on 3 format	Generati on 4 format	Generati on 5 format	Generati on 5A format	Generati on 6 format	Generati on 7 format
Generatio n 1	Read/ write access	n/a	n/a	n/a	n/a	n/a	n/a	n/a
Generatio n 2	Read/ write access	Read/ write access	n/a	n/a	n/a	n/a	n/a	n/a
Generatio n 3	Read- only access	Read/ write access	Read/ write access	n/a	n/a	n/a	n/a	n/a
Generatio n 4	Read- only access	Read- only access	Read/ write access	Read/ write access	n/a	n/a	n/a	n/a
Generatio n 5	n/a	n/a	n/a	Read/ write access	Read/ write access	n/a	n/a	n/a
Generatio n 5A	n/a	n/a	n/a	Read- only access	Read/ write access	Read/ write access	n/a	n/a

Drives	Generati on 1 format	Generati on 2 format	Generati on 3 format	Generati on 4 format	Generati on 5 format	Generati on 5A format	Generati on 6 format	Generati on 7 format
Generatio n 6	n/a	n/a	n/a	Read/ write access	Read/ write access	Read/ write access	Read/ write access	n/a
Generatio n 7	n/a	n/a	n/a	n/a	n/a	n/a	n/a	Read/ write access

If you must mix generations of drives in a library, review the example and restrictions to help prevent problems.

Table 23. Mixing generations of drives

Library type	Example and restrictions
SCSI	<p>Define a new storage pool and device class for the latest drive generation. For example, suppose that you have a storage pool and device class for 3592-2. The storage pool contains all the media that were written in Generation 2 format. Suppose that the value of the FORMAT parameter in the device class definition is set to 3592-2 (not DRIVE). You add Generation 3 drives to the library. Complete the following steps:</p> <ol style="list-style-type: none"> 1. In the new device-class definition for the Generation 3 drives, set the value of the FORMAT parameter to 3592-3 or 3592-3C. Do not specify DRIVE. 2. In the definition of the storage pool that is associated with Generation 2 drives, update the MAXSCRATCH parameter to 0, for example: <pre>update stgpool genpool2 maxscratch=0</pre> <p>This method allows both generations to use their optimal format and minimizes potential media problems that can result from mixing generations. However, it does not resolve all media issues. For example, competition for mount points and mount failures might result. (To learn more about mount point competition in the context of 3592 drives and media, see “Defining 3592 device classes” on page 97.)</p> <p>Restriction: The following list describes media restrictions:</p> <ul style="list-style-type: none"> • CHECKIN LIBVOL: The issue is using the CHECKLABEL=YES option. If the label is written in a Generation 3 or later format, and you specify the CHECKLABEL=YES option, drives of previous generations fail by using this command. To avoid the issue, specify CHECKLABEL=BARCODE. • LABEL LIBVOL: When the server tries to use drives of a previous generation to read the label that is written in a Generation 3 or later format, the LABEL LIBVOL command fails unless OVERWRITE=YES is specified. Verify that the media that is being labeled with OVERWRITE=YES does not have any active data. • CHECKOUT LIBVOL: When IBM Storage Protect verifies the label (CHECKLABEL=YES) as a Generation 3 or later format, and read drives of previous generations, the command fails. To avoid this issue, specify CHECKLABEL=NO.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[CHECKOUT LIBVOLUME \(Check a storage volume out of a library\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

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

Controlling data-access speeds for 3592 volumes

You can optimize the storage capacity and improve data-access speeds when you create volumes. By partitioning data into storage pools that have volumes, you can specify the scale capacity percentage to provide maximum storage capacity, or to provide fast access to the volume.

About this task

To reduce media capacity, specify the **SCALECAPACITY** parameter when you define the device class by using the **DEFINE DEVCLASS** command or when you update the device class by using the **UPDATE DEVCLASS** command.

Specify a percentage value of 20, 90, or 100. A value of 20 percent provides the fastest access time, and 100 percent provides the largest storage capacity. For example, if you specify a scale capacity of 20 for a 3592 device class without compression, a 3592 volume in that device class would store 20 percent of its full capacity of 300 GB, or about 60 GB.

Scale capacity takes effect only when data is first written to a volume. Updates to the device class for scale capacity do not affect volumes that already have data written to them until the volume is returned to scratch status.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Enabling and disabling 3592 Generation 2 and later drive encryption

With IBM Storage Protect, you can use the following types of drive encryption with drives that are 3592 Generation 2 and later: Application, System, and Library. These methods are defined through the hardware.

About this task

The **DRIVEENCRYPTION** parameter on the **DEFINE DEVCLASS** command specifies whether drive encryption is allowed for drives that are 3592 Generation 2 and later. Use this parameter to ensure IBM Storage Protect compatibility with hardware encryption settings for empty volumes. You cannot use this parameter for storage pool volumes that are full or are filling.

- To use the Application method, in which IBM Storage Protect generates and manages encryption keys, set the **DRIVEENCRYPTION** parameter to **ON**. This enables the encryption of data for empty volumes. If the parameter is set to **ON** and if the hardware is configured for another encryption method, backup operations fail.
- To use the Library or System methods of encryption, set the parameter to **ALLOW**. This specifies that IBM Storage Protect is not the key manager for drive encryption, but allows the hardware to encrypt the volume's data through one of the other methods. Specifying this parameter does not automatically encrypt volumes. Data can be encrypted only by specifying the **ALLOW** parameter and configuring the hardware to use one of these methods.

The **DRIVEENCRYPTION** parameter is optional. The default value is to allow the Library or System methods of encryption.

Procedure

The following simplified example shows how to encrypt data for empty volumes in a storage pool, by using IBM Storage Protect as the key manager:

1. Define a library by issuing the **DEFINE LIBRARY** command.
For example, issue the following command:

```
define library 3584 libtype=SCSI
```

2. Define a device class, 3592_ENCRYPT, by issuing the **DEFINE DEVCLASS** command and specifying the value ON for the **DRIVEENCRYPTION** parameter.
For example, issue the following command:

```
define devclass 3592_encrypt library=3584 devtype=3592 driveencryption=on
```

3. Define a storage pool.
For example, issue the following command:

```
define stgpool 3592_encrypt_pool 3592_encrypt
```

What to do next

To disable any method of encryption on new volumes, set the **DRIVEENCRYPTION** parameter to OFF. If the hardware is configured to encrypt data through either the Library or System method and **DRIVEENCRYPTION** is set to OFF, backup operations fail.

Configuring library sharing

Multiple IBM Storage Protect servers can share storage devices by using a storage area network (SAN). You set up one server as the library manager and the other servers as library clients.

Before you begin

Ensure that your systems meet licensing requirements for library sharing. An entitlement for IBM Storage Protect for SAN is required for each IBM Storage Protect server that is configured as a library client or a library manager in a SAN environment.

About this task

With LAN-free data movement, IBM Storage Protect client systems can directly access storage devices that are defined to an IBM Storage Protect server. Storage agents are installed and configured on the client systems to perform the data movement.

To set up library sharing, you must define one IBM Storage Protect server as the library manager for your shared library configuration. Then, you define other IBM Storage Protect servers as library clients that communicate and request storage resources from the library manager. The library manager server must be at the same version or a later version as the server or servers that are defined as library clients.

Procedure

To complete the following steps to share library resources on a SAN among IBM Storage Protect servers, complete the following steps:

1. Set up server-to-server communications.

To share a storage device on a SAN, define servers to each other by using the cross-define function. Each server must have a unique name.

2. Define a shared library and set up tape devices on the server systems.

Use the procedure that is described in [“Configuring libraries for use by a server” on page 90](#) to define a library for use in the shared environment. Modify the procedure to define the library as shared, by specifying the **SHARED=YES** parameter for the **DEFINE LIBRARY** command.

3. Define the library manager server.
4. Define the shared library on the server that is the library client.
5. From the library manager server, define paths from the library client to each drive that the library client can access.

The device name must reflect the way that the library client system recognizes the tape device. A path from the library manager to each tape drive must be defined in order for the library client to use the drive.

To avoid problems, ensure that all drive path definitions that are defined for the library manager are also defined for each library client.

For example, if the library manager defines three tape drives, the library client must also define three tape drives. To limit the number of tape drives that a library client can use at a time, use the **MOUNTLIMIT** parameter of the device class on the library client.

6. Define device classes for the shared library.

The preferred method is to make the device class names the same on both servers to avoid confusion when you define multiple device classes with the same device type and library parameters. Some operations, such as database backup, use the device class name to identify the data for backup.

The device class parameters that are specified on the library manager override the parameters that are specified for the library client. If the device class names are different, the library manager uses the parameters that are specified in a device class that matches the device type that is specified for the library client.

7. Define a storage pool for the shared library.

8. Repeat the steps to configure another server as a library client.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE STGPPOOL \(Define a volume in a storage pool\)](#)

Linux

AIX

Example: Library sharing for AIX and Linux servers

To learn how to set up a SCSI library sharing environment for servers that run on AIX or Linux systems, review the sample procedure.

About this task

In this example, a library manager server named ASTRO and a library client named JUDY are configured. To help clarify where each step is performed, the commands are preceded by the server name from which the command is issued. Most commands are issued from the library client.

For SCSI libraries, define the library by specifying the **libtype=scsi** parameter.

Procedure

1. To set up ASTRO as the library manager server, define a shared SCSI library named SANGROUP.

For example:

```
astro> define library sangroup libtype=scsi shared=yes
```

Then complete the rest of the steps as described in [Example: Configure a SCSI or virtual tape library with a single drive device type](#) to configure the library.

Tip: You can use the **PERFORM LIBACTION** command to define drives and paths for a library in one step.

2. Define ASTRO as the library manager server by issuing the **DEFINE SERVER** command.

```
judy> define server astro serverpassword=secret hladdress=192.0.2.24  
lladdress=1777 crossdefine=yes
```

3. Define the shared library SANGROUP by issuing the **DEFINE LIBRARY** command. You must use the library manager server name in the **PRIMARYLIBMANAGER** parameter, and use **LIBTYPE=SHARED**.

```
judy> define library sangroup libtype=shared primarylibmanager=astro
```

Ensure that the library name is the same as the library name on the library manager.

4. Define paths from the library manager, ASTRO, to two drives in the shared library by issuing the **DEFINE PATH** command.

```
AIX astro> define path judy drivea srctype=server desttype=drive
library=sangroup device=/dev/rmt6
astro> define path judy driveb srctype=server desttype=drive
library=sangroup device=/dev/rmt7
```

```
Linux astro> define path judy drivea srctype=server desttype=drive
library=sangroup device=/dev/IBMtape6
astro> define path judy driveb srctype=server desttype=drive
library=sangroup device=/dev/IBMtape7
```

5. Define all device classes that are associated with the shared library.

```
AIX judy> define devclass tape library=sangroup devtype=lto
```

```
Linux judy> define devclass tape library=sangroup devtype=lto
```

The following parameters for the device class definition must be the same on the library client and the library manager:

- **LIBRARY**
- **DRIVEENCRYPTION**
- **WORM**
- **FORMAT**

6. Define a storage pool that is named BACKTAPE for the shared library to use. Issue the **DEFINE STGPOOL** command.

```
judy> define stgpool backtape tape maxscratch=50
```

What to do next

Repeat the procedure to define more library clients to your library manager.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE DRIVE \(Define a drive to a library\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Windows Example: Library sharing for Windows servers

To learn how to set up a library sharing environment for servers that run on Windows systems, review the sample procedure.

About this task

In this example, a library manager server named ASTRO and a library client named JUDY are configured.

For SCSI libraries, define the library by specifying the **libtype=scsi** parameter.

Windows Setting up the library manager server

You must set up the library manager server in order to configure the IBM Storage Protect servers to share SAN-connected devices.

Procedure

The following procedure is an example of how to set up an IBM Storage Protect server that is named ASTRO as a library manager:

1. Ensure that the library manager server is running:
 - a) Start the Windows Services Management Console (services.msc).
 - b) Select the service. For example, TSM Server1.
 - c) If the service is not running, right-click the service name and click **Start**.
2. Obtain the library and drive information for the shared library device:
 - a) Run the `tsmdlst.exe` utility. The utility is in the `\Program Files\Tivoli\TSM\server` directory.
3. Define a library whose library type is SCSI.

For example:

```
define library sangroup libtype=scsi shared=yes
```

This example uses the default for the library's serial number, which is to have the server obtain the serial number from the library itself at the time that the path is defined. Depending on the capabilities of the library, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device.

4. Define the path from the server to the library.

```
define path astro sangroup srctype=server desttype=library  
device=lb0.0.0.2
```

If you did not include the serial number when you defined the library, the server now queries the library to obtain this information. If you did include the serial number when you defined the library, the server verifies what you defined and issues a message if there is a mismatch.

5. Define the drives in the library.

```
define drive sangroup drivea  
define drive sangroup driveb
```

This example uses the default for the drive's serial number, which is to have the server obtain the serial number from the drive itself at the time that the path is defined. Depending on the capabilities of the drive, the server might not be able to automatically detect the serial number. In this case, the server does not record a serial number for the device, and is not able to confirm the identity of the device when you define the path or when the server uses the device.

This example also uses the default for the drive's element address, which is to have the server obtain the element number from the drive itself at the time that the path is defined.

The element address is a number that indicates the physical location of a drive within an automated library. The server needs the element address to connect the physical location of the drive to the drive's SCSI address. You can have the server obtain the element number from the drive itself at the time that the path is defined, or you can specify the element number when you define the drive.

Depending on the capabilities of the library, the server might not be able to automatically detect the element address. In this case, you must supply the element address when you define the drive. Element numbers for many libraries are available at [IBM Support Portal for IBM Storage Protect](#).

6. Define the path from the server to each of the drives.

```
define path astro drivea srctype=server desttype=drive library=sangroup
device=mt0.1.0.2
define path astro driveb srctype=server desttype=drive library=sangroup
device=mt0.2.0.2
```

If you did not include the serial number or element address when you defined the drive, the server now queries the drive or the library to obtain this information.

7. Define at least one device class.

```
define devclass tape devtype=dlt library=sangroup
```

8. Check in the library inventory. The following example checks all volumes into the library inventory as scratch volumes. The server uses the name on the bar code label as the volume name.

```
checkin libvolume sangroup search=yes status=scratch
checklabel=barcode
```

9. Set up a storage pool for the shared library with a maximum of 50 scratch volumes.

```
define stgpool backtape tape
description='storage pool for shared sangroup' maxscratch=50
```

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE DRIVE \(Define a drive to a library\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Windows Setting up the library client servers

You must set up one or more library client servers to configure the IBM Storage Protect servers to share SAN-connected devices.

Before you begin

Ensure that a library manager server is defined.

About this task

You must define the library manager server. Use the following procedure as an example of how to set up an IBM Storage Protect server that is named JUDY as a library client.

Procedure

1. Ensure that the library manager server is running:
 - a) Start the Windows Services Management Console (services.msc).
 - b) Select the service. For example, TSM Server1.
 - c) If the service is not running, right-click and select **Start**.
2. Obtain the library and drive information for the shared library device:
 - a) Run the `tsmdl1st.exe` utility. The utility is in the `\Program Files\Tivoli\TSM\server` directory.
3. Define the shared library, SANGROUP, and identify the library manager. Ensure that the library name is the same as the library name on the library manager.

```
define library sangroup libtype=shared primarylibmanager=astro
```

4. Define the paths from the library client server to each of the drives by issuing commands on the administrative client:

```
define path judy drivea srctype=server desttype=drive library=sangroup
device=mt0.1.0.3
define path judy driveb srctype=server desttype=drive library=sangroup
device=mt0.2.0.3
```

5. Define at least one device class by issuing commands from the library client:

```
define devclass tape devtype=dlt mountretention=1 mountwait=10
library=sangroup
```

Set the parameters for the device class the same on the library client as on the library manager. Making the device class names the same on both servers is also a good practice, but is not required.

The device class parameters that are specified on the library manager server override those specified for the library client. This is true whether or not the device class names are the same on both servers. If the device class names are different, the library manager uses the parameters specified in a device class that matches the device type specified for the library client.

If a library client requires a setting that is different from what is specified in the library manager's device class (for example, a different mount limit), complete the following steps:

- a. Create an additional device class on the library manager server. Specify the parameter settings that you want the library client to use.
 - b. Create a device class on the library client with the same name and device type as the new device class you created on the library server.
6. Define the storage pool, BACKTAPE, that will use the shared library:

```
define stgpool backtape tape
description='storage pool for shared sangroup' maxscratch=50
```

7. Repeat this procedure to define additional servers as library clients.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE PATH \(Define a path\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Setting up a storage pool hierarchy

As part of the implementation process, you must set up a storage pool hierarchy. Set up at least one primary storage pool on disk and one primary storage pool on tape. Ensure that data is migrated from disk to tape daily.

Before you begin

1. Ensure that you reviewed the information in [“Planning the storage pool hierarchy” on page 21](#).
2. Ensure that appropriate rules, also known as *policies*, are specified for backing up client data. Follow the instructions in [“Specifying rules for backing up and archiving client data” on page 109](#).
3. Ensure that a policy is assigned to each node. For instructions about assigning a policy when you register a node, see [“Registering clients” on page 113](#).

Procedure

To set up a storage pool hierarchy, complete the following steps:

1. Define a primary storage pool for the tape device by issuing the **DEFINE STGPOOL** command.

For example, define a primary storage pool, TAPE1, with a device class of LTO, and enable group collocation. Set the maximum number of scratch volumes that the server can request for this storage pool to 999. Issue the following command:

```
define stgpool tape1 lto pooltype=primary collocate=group  
maxscratch=999
```

2. Define the drives, paths, and libraries for the primary storage pool on tape. Follow the instructions in [“Defining tape devices”](#) on page 91.
3. Define a primary storage pool for the disk device by issuing the **DEFINE STGPOOL** command.

For example, define a storage pool, DISK1, with a device class of FILE. Ensure that data can be migrated to the tape storage pool, TAPE1, but prevent automatic migration by specifying 100 for the **HIGHMIG** parameter and 0 for the **LOWMIG** parameter. Prevent reclamation by specifying 100 for the **RECLAIM** parameter. Enable node collocation. Set the maximum number of scratch volumes that the server can request for this storage pool to 9999. Use the **MIGPROCESS** parameter to specify the number of migration processes. The value of the **MIGPROCESS** parameter should equal the number of drives in the library minus the number of drives that are reserved for restore operations. Issue the following command:

```
define stgpool disk1 file pooltype=primary nextstgpool=tape1  
highmig=100 lowmig=0 reclaim=100 collocate=node maxscratch=9999 migprocess=5
```

For more information about how to set up migration from disk to tape, see *Planning the storage pool hierarchy* in IBM Documentation.

What to do next

A storage pool hierarchy includes only primary storage pools. After you set up the storage pool hierarchy, complete the following steps:

1. Create a copy storage pool on a tape device. For instructions, see [DEFINE STGPOOL \(Define a copy storage pool assigned to sequential access devices\)](#).
2. Back up the tape-based primary storage pool to the copy storage pool by using the **BACKUP STGPOOL** command. For instructions, see [BACKUP STGPOOL \(Back up primary storage pool data to a copy storage pool\)](#).
3. To ensure that data can be recovered in a disaster, set up a procedure for moving tape volumes from the copy storage pool to an offsite location. For instructions, see [“Preparing for and recovering from a disaster by using DRM”](#) on page 208.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Protecting applications and systems

The server protects data for clients, which can include applications, virtual machines, and systems.

Adding clients

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

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 108.
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 109.
3. Specify when to back up and archive client data. Follow the instructions in [“Scheduling backup and archive operations”](#) on page 112.
4. To allow the client to connect to the server, register the client. Follow the instructions in [“Registering clients”](#) on page 113.
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 114.

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 114](#) 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
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)

Goal	Product and description	Installation instructions
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 112.
- 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 110.
2. If you need to make minor changes to meet data retention requirements, follow the instructions in [“Editing policies”](#) on page 111.
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>
Modify a management class	<p>To change the properties of a management class, update the fields in the table.</p>

6. Click **Save**.



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

7. Click **Activate**. A summary of the differences between the active policy set and the new policy set is displayed. Ensure that the changes in the new policy set are consistent with your data retention requirements by completing the following steps:
 - a) Review the differences between corresponding management classes in the two policy sets, and consider the consequences for client files. Client files that are bound to management classes in the active policy set will be bound to the management classes with the same names in the new policy set.
 - b) Identify management classes in the active policy set that do not have counterparts in the new policy set, and consider the consequences for client files. Client files that are bound to these management classes will be managed by the default management class in the new policy set.
 - c) If the changes to be implemented by the policy set are acceptable, select the **I understand that these updates can cause data loss** check box and click **Activate**.

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.

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 113](#).

- 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

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

Option	Description
Modify a schedule	<ol style="list-style-type: none"> a. In the Schedules view, select the schedule and click Details. 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 115](#).

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

Software	Link to instructions
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 113](#)).
 - 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. Optional: 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](#).
4. Configure the client to run scheduled operations. Follow the instructions in [“Configuring the client to run scheduled operations” on page 116](#).
5. Optional: Configure communications through a firewall. Follow the instructions in [“Configuring client/server communications through a firewall” on page 118](#).
6. Run a test backup to verify that data is protected as you planned.
For example, for a backup-archive client, complete the following steps:

- a) On the **Clients** page of the Operations Center, select the client that you want to back up, and click **Back Up**.
 - b) Verify that the backup completes successfully and that there are no warning or error messages.
7. Monitor the results of the scheduled operations for the client in the Operations Center.

What to do next

If you need to change what is getting backed up from the client, follow the instructions in [“Modifying the scope of a client backup” on page 112](#).

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

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

```
dsmc query sess
```

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

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

- g. To enable the client acceptor to start automatically after a system restart, add the following entry to the system startup file (typically, `/etc/inittab`):

```
tсм::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
```

- f. 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.

- i. 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.

Configuring LAN-free data movement

You can configure the client and server so that the client, through a storage agent, can move data directly to storage on a SAN.

About this task

LAN-free data movement is provided by the IBM Storage Protect for SAN product. For more information, see the documentation for [IBM Storage Protect for SAN](#).

Procedure

To configure LAN-free data movement, complete the following steps.

1. Verify the network connection.
2. Establish communications among the client, storage agent, and the server.
3. Install and configure software on client systems.
4. Configure devices on the server for the storage agent to access.
5. Configure IBM Storage Protect policies for LAN-free data movement for the client.
6. If you are using shared FILE storage, install and configure IBM Storage Scale.

Restriction: **Windows** If an IBM Storage Scale volume is formatted by an AIX server, the Windows system uses TCP/IP to transfer data and not the storage area network.

7. Define paths from the storage agent to drives.
8. Start the storage agent and verify the LAN-free configuration.

What to do next

To help you tune the use of your LAN and SAN resources, you can control the path that data transfers take for clients with the capability of LAN-free data movement. To control the path, run the following command: **UPDATE NODE**

For each client, you can select one of the following settings for data read and write operations. Specify data read operations by using the **DATAREADPATH** parameter and data write operations by using the **DATAWRITEPATH** parameter. The parameter is optional. The default value is ANY.

LAN (LAN path only)

If either of the following conditions is true, specify the LAN value:

- You want to back up or restore a small amount of data.
- The client does not have SAN connectivity.

LANFREE (LAN-free path only)

If the client and server are on the same SAN and any of the following conditions are true, specify the LANFREE value:

- You want to back up or restore a large amount of data.
- You want to offload the server processing load to the client.
- You want to relieve LAN congestion.

ANY (Any available path)

If a LAN-free path is available that path is used. If a LAN-free path is not available, the data is moved by using the LAN.

Validating your LAN-free configuration

After you configure an IBM Storage Protect client for LAN-free data movement, you can verify the configuration and server definitions by using the **VALIDATE LANFREE** command.

About this task

The **VALIDATE LANFREE** command allows you to determine which destinations for a node that is using a specific storage agent are capable of LAN-free data movement. The command output can also help identify if there is a problem with an existing LAN-free configuration. You can evaluate the policy, storage pool, and path definitions for a node and storage agent that the node is using to ensure that an operation is working properly.

Procedure

- Determine whether a client node has a problem with its LAN-free configuration by issuing the **VALIDATE LANFREE** command. For example, if the client node FRED is using the storage agent FRED_STA, issue the following command:

```
validate lanfree fred fred_sta
```

The results help you to identify adjustments that might be needed in the storage configuration or policies. The output displays which management class destinations for a specific operation type are not capable of LAN-free data transfers. It also reports the total number of LAN-free destinations.

Related information

[VALIDATE LANFREE \(Validate LAN-Free paths\)](#)

Tape encryption methods

Deciding on the encryption method to use depends on how you want to manage your data.

It is critical to secure client data, especially when that data is sensitive. To ensure that data in onsite and offsite volumes is protected, IBM tape encryption technology is available.

This technology uses a stronger level of encryption by requiring 256-bit Advanced Encryption Standard (AES) encryption keys. Keys are passed to the drive by a key manager to encrypt and decrypt data.

IBM tape technology supports different methods of drive encryption for the following devices:

- IBM 3592 Generation 2 and later generations
- IBM Linear Tape-Open Generation 4 and later generations

The methods of drive encryption that you can use with IBM Storage Protect are set up at the hardware level. IBM Storage Protect cannot control or change which encryption method is used in the hardware configuration. If the hardware is set up for the Application method, IBM Storage Protect can turn encryption on or off depending on the **DRIVEENCRYPTION** value on the device class.

To encrypt all data in a particular logical library or to encrypt data on more than just storage pool volumes, use the Library or System method. If the encryption key manager is set up to share keys, the Library and System methods can share the encryption key, which allows the two methods to be interchanged. IBM Storage Protect cannot share or use encryption keys between the Application method and either the Library or the System methods of encryption.

Table 24. Encryption methods

Encryption method	Description
Application encryption	<p>With application-managed encryption, you can create dedicated storage pools that contain encrypted volumes only. This way, you can use storage pool hierarchies and policies to manage the way data is encrypted.</p> <p>Encryption keys are managed by the application, in this case, IBM Storage Protect. IBM Storage Protect generates and stores the keys in the server database. Data is encrypted during write operations, when the encryption key is passed from the server to the drive. Data is decrypted for read operations.</p> <p>To encrypt storage pool volumes and eliminate some encryption processing on your system, enable the Application method. Use application-managed encryption only for storage pool volumes. Other volumes, such as backup-set tapes, export volumes, and database backups, are not encrypted by using the Application method.</p> <p>Requirement: When application encryption is enabled, you must take extra care to secure database backups because the encryption keys that are used to encrypt and decrypt data are stored in the server database. To restore your data, you must have the correct database backup and corresponding encryption keys to access your information. Ensure that you back up the database frequently and safeguard the backups to prevent data loss or theft. Anyone who has access to both the database backup and the encryption keys has access to your data.</p>
Library encryption	<p>With library-managed encryption, you can control which volumes are encrypted by using their serial numbers. You can specify a range or set of volumes to encrypt.</p> <p>Encryption keys are managed by the library. Keys are stored in an encryption key manager and provided to the drive. If you set up the hardware to use library-managed encryption, you can use this method by running the DEFINE DEVCLASS command and specifying the DRIVEENCRYPTION=ALLOW parameter.</p> <p>Restriction: Only certain IBM libraries support IBM LTO-4 and later encryption. For more information, see “Configuring tape drive encryption” on page 122.</p>

Table 24. Encryption methods (continued)	
Encryption method	Description
System encryption	System-managed encryption is available only on the AIX® operating system. Encryption keys that are provided to the drive are managed by the device driver or operating system and stored in an encryption key manager. If the hardware is set up to use system encryption, you can use this method by running the DEFINE DEVCLASS command and specifying the DRIVEENCRYPTION=ALLOW parameter.

To determine whether a volume is encrypted and which method was used, run the **QUERY VOLUME** command and specify the **FORMAT=DETAILED** parameter.

Configuring tape drive encryption

You can use drive encryption to protect tapes that contain critical or sensitive data, for example, tapes that contain confidential financial information. Drive encryption can be useful when you move tapes from the IBM Storage Protect server environment to an onsite or offsite location.

About this task

To determine which encryption methods can be used with various drive types, see the following table.

Table 25. Available encryption methods			
	Application method	Library method	System method
3592 Generation 2 and later	Yes	Yes.	Yes
HP LTO-4 and later	Yes	No.	No
IBM LTO-4 and later	Yes	Yes, but only if your system hardware (for example, a TS3500 tape library) supports it.	Yes
Oracle StorageTek T10000B	Yes	No.	No
Oracle StorageTek T10000C	Yes	No.	No
Oracle StorageTek T10000D	Yes	No.	No

A library can contain a mixture of drives, some of which support encryption and some of which do not. For example, a library might contain two LTO-2 drives, two LTO-3 drives, and two LTO-4 drives. You can also mix media in a library by using, for example, encrypted and non-encrypted device classes that have different tape and drive technologies.

Restrictions:

- To apply encryption to LTO-4 or later drives, all of the drives must support encryption.
- To apply encryption to a logical library, you must use the same method of encryption for all drives within the library. Do not create an environment in which some drives use the Application method and some drives use the Library or System methods of encryption.

For more information about setting up your hardware environment to use drive encryption, see your hardware documentation.

Procedure

1. Install a device driver that supports drive encryption:
 - To enable encryption for an IBM LTO-4 or later drive, you must install the IBM RMSS Ultrium device driver. SCSI drives do not support IBM LTO-4 or later encryption.
 - To enable encryption for an HP LTO-4 or later drive, you must install the IBM Storage Protect device driver.
2. Enable drive encryption by specifying the **DRIVEENCRYPTION** parameter on the **DEFINE DEVCLASS** or **UPDATE DEVCLASS** command for the 3592, LTO, or ECARTRIDGE device types.

What to do next

When you use encryption-capable drives with a supported encryption method, a different format is used to write encrypted data to tapes. When data is written to volumes that use the different format and if the volumes are then returned to scratch, they contain labels that can be read only by encryption-enabled drives. To use these scratch volumes in a drive that is not enabled for encryption, either because the hardware is not capable of encryption or because the encryption method is set to NONE, you must relabel the volumes.

Related tasks

[Enabling and disabling 3592 Generation 2 and later drive encryption](#)

With IBM Storage Protect, you can use the following types of drive encryption with drives that are 3592 Generation 2 and later: Application, System, and Library. These methods are defined through the hardware.

[Enabling and disabling drive encryption for LTO Generation 4 or later tape drives](#)

IBM Storage Protect supports the three types of drive encryption that are available with LTO Generation 4 or later drives: Application, System, and Library. These methods are defined through the hardware.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Controlling tape storage operations

Device class definitions for tapes include parameters that allow you to control storage operations.

How IBM Storage Protect fills volumes

The **DEFINE DEVCLASS** command has an optional **ESTCAPACITY** parameter that indicates the estimated capacity for sequential volumes that are associated with the device class. IBM Storage Protect uses the estimated capacity of volumes to determine the estimated capacity of a storage pool, and the estimated percent utilized.

If the **ESTCAPACITY** parameter is not specified, IBM Storage Protect uses a default value that is based on the recording format that is specified for the device class by using the **FORMAT** parameter.

If you specify an estimated capacity that exceeds the actual capacity of the volume in the device class, IBM Storage Protect updates the estimated capacity of the volume when the volume becomes full. When IBM Storage Protect reaches the end of the volume, it updates the capacity to match the amount that is written to the volume.

You can either accept the default estimated capacity for the device class, or explicitly specify an estimated capacity. An accurate estimated capacity value is not required, but is useful. IBM Storage Protect uses the estimated capacity of volumes to determine the estimated capacity of a storage pool,

and the estimated percent that is used. You might want to change the estimated capacity if on or both of the following conditions are true:

- The default estimated capacity is inaccurate because of data compression.
- You have volumes of nonstandard size.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Specifying the estimated capacity of tape volumes

IBM Storage Protect also uses estimated capacity to determine when to begin the reclamation of storage pool volumes.

About this task

For tape device classes, the default values selected by the server depend on the recording format that is used to write data to the volume. You can either accept the default for a device type or specify a value.

To specify estimated capacity for tape volumes, use the **ESTCAPACITY** parameter when you define the device class or update its definition.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Specifying recording formats for tape media

You can specify the recording format that is used by IBM Storage Protect to write data to tape media. If you plan to mix generations of drives, or different drive types, within a library, you must specify a recording format for each drive generation and each drive type. In this way, the server can differentiate between the drive generations and drive types.

About this task

To specify a recording format, use the **FORMAT** parameter when you define the device class or update its definition.

If all drives associated with that device class are identical, specify **FORMAT=DRIVE**. The server selects the highest format that is supported by the drive on which a volume is mounted.

If some drives associated with the device class support a higher density format than others, specify a format that is compatible with all drives.

If drives in a single SCSI library use different tape technologies (for example, DLT and LTO Ultrium), specify a unique value for the **FORMAT** parameter in each device class definition.

For a configuration example, see [Example: Configure a SCSI or virtual tape library with multiple drive device types](#).

The recording format that the server uses for a volume is selected when data is first written to the volume. Updating the **FORMAT** parameter does not affect media that already contain data until those media are rewritten from the beginning. This process might happen after a volume is reclaimed or deleted, or after all of the data on the volume expires.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Associating library objects with device classes

A library contains the drives that can be used to mount the volume. Only one library can be associated with a device class. However, multiple device classes can reference the same library.

About this task

To associate a device class with a library, use the **LIBRARY** parameter when you define a device class or update its definition.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Controlling media-mount operations for tape devices

By using device class definitions, you can control the number of mounted volumes, the amount of time a volume remains mounted, and the amount of time that the IBM Storage Protect server waits for a drive to become available.

Controlling the number of simultaneously mounted volumes

When you set a mount limit for a device class, you must consider the number of storage devices that are connected to your system. You must also consider whether you use the simultaneous-write function, whether you associate multiple device classes with a single library, and the number of processes that run at the same time.

About this task

When you select a mount limit for a device class, consider the following issues:

- How many storage devices are connected to your system?

Do not specify a mount limit value that is greater than the number of associated available drives in your installation. If the server tries to mount as many volumes as specified by the mount limit and no drives are available for the required volume, an error occurs and client sessions might end. (This restriction does not apply when the **DRIVES** parameter is specified.)

If you are sharing library resources on a SAN among IBM Storage Protect servers, you must limit the number of tape drives that a library client can use at a time. To allow multiple library client servers use a library simultaneously specify the **MOUNTLIMIT** parameter when you define or update the device class on the library client. For more information about configuring library sharing, see [“Configuring library sharing” on page 100](#).

- Are you using the simultaneous-write function to primary storage pools, copy storage pools, and active-data pools?

Specify a mount limit value that provides enough mount points to support writing data simultaneously to the primary storage pool and all associated copy storage pools and active-data pools.

- Are you associating multiple device classes with a single library?

A device class that is associated with a library can use any drive in the library that is compatible with the device class' device type. Because you can associate more than one device class with a library, a single drive in the library can be used by more than one device class. IBM Storage Protect ensures that two operations cannot use the same drive simultaneously by using two different device classes.

- How many IBM Storage Protect processes do you want to run at the same time, by using devices in this device class?

IBM Storage Protect automatically cancels some processes to run other, higher priority processes. If the server is using all available drives in a device class to complete higher priority processes, lower-priority processes must wait until a drive becomes available. For example, IBM Storage Protect cancels the

process for a client that backs up directly to tape if the drive is needed for a server migration or tape reclamation process. IBM Storage Protect cancels a tape reclamation process if the drive is needed for a client restore operation. For more information, see [“Preempting operations” on page 127](#).

If processes are often canceled by other processes, consider whether you can make more drives available for IBM Storage Protect use. Otherwise, review your scheduling of operations to reduce the contention for drives.

This consideration also applies to the simultaneous-write function. You must have enough drives available to allow for a successful simultaneous-write operation.

To specify the maximum number of volumes that can be simultaneously mounted, use the **MOUNTLIMIT** parameter when you define the device class or update its definition.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Controlling the amount of time that a volume remains mounted

You can control the amount of time that a mounted volume remains mounted after its last I/O activity. If a volume is used frequently, you can improve performance by setting a longer mount retention period to avoid unnecessary mount and dismount operations.

About this task

If mount operations are being handled by manual, operator-assisted activities, you might want to specify a long mount retention period. For example, if only one operator supports your entire operation on a weekend, then define a long mount retention period so that the operator is not being asked to mount volumes every few minutes.

To control the amount of time a mounted volume remains mounted, use the **MOUNTRETENTION** parameter when you define the device class or update its definition. For example, if the mount retention value is 60, and a mounted volume remains idle for 60 minutes, the server dismounts the volume.

While IBM Storage Protect has a volume mounted, the drive is allocated to IBM Storage Protect and cannot be used for anything else. If you need to free the drive for other uses, you can cancel IBM Storage Protect operations that are using the drive and then dismount the volume. For example, you can cancel server migration or backup operations. For information on how to cancel processes and dismount volumes, see [“Managing server requests for volumes” on page 189](#)

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Controlling the amount of time that the server waits for a drive

You can specify the maximum amount of time, in minutes, that the IBM Storage Protect server waits for a drive to become available for the current mount request.

About this task

To control the wait time for a drive to become available for a mount request, use the **MOUNTWAIT** parameter when you define or update a device class.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Preempting operations

The server can preempt server or client operations for a higher priority operation when a mount point is in use and no others are available, or access to a specific volume is required. When an operation is preempted, it is canceled.

You can use the **QUERY MOUNT** command to see the status of the volume for the mount point.

By default, preemption is enabled on the server. To disable preemption, specify the **NOPREEMPT** option in the server options file. If you specify this option, the **BACKUP DB** command, and the export and import commands are the only operations that can preempt other operations.

Related information

[BACKUP DB \(Back up the database\)](#)

[QUERY MOUNT \(Display information on mounted sequential access volumes\)](#)

Mount point preemption

If a high-priority operation requires a mount point that is in a specific device class and all the mount points in the device class are in use, the high-priority operation can preempt a mount point from a lower-priority operation.

Mount points can be preempted only when the device class of the operation preempting and the operation that is being preempted is the same.

The following high-priority operations can preempt other operations for a mount point.

- Database backup operations
- Retrieve, restore, or HSM recall operations that are initiated by clients
- Restore operations by using a remote data mover
- Export operations
- Import operations
- Operations to generate backup sets

The following server operations cannot preempt other operations or be preempted:

- Audit a volume
- Restore data from a copy or active-data pool
- Prepare a recovery plan file
- Store data by using a remote data mover

The following operations can be preempted and are listed in order of priority, from highest priority to lowest priority. The server selects the lowest priority operation to preempt, for example, identify duplicates.

- Replicate nodes
- Back up data to a copy storage pool
- Copy active data to an active data pool
- Move data on a storage pool volume
- Migrate data from disk to sequential media
- Migrate data from sequential media to sequential media
- Back up, archive, or HSM migrate operations that are initiated by clients
- Reclaim volumes in a sequential-access storage pool
- Identify duplicates

Volume access preemption

If a high-priority operation requires access to a specific volume and that volume is in use, the high-priority operation can preempt the lower-priority operation for that volume.

For example, if a restore request requires access to a volume in use by a reclamation operation and a drive is available, the reclamation operation is canceled.

The following high-priority operations can preempt operations for access to a specific volume:

- Database backup operations
- Retrieve, restore, or HSM recall operations that are initiated by clients
- Restore operations by using a remote data mover
- Export operations
- Import operations
- Operations to generate backup sets

The following operations cannot preempt other operations or be preempted:

- Audit volume
- Restore data from a copy or active-data pool
- Prepare a recovery plan
- Store data by using a remote data mover

The following operations can be preempted, and are listed in order of priority, from highest priority to lowest priority. The server selects the lowest priority operation to preempt, for example, identify duplicates.

- Replicate nodes
- Back up data to a copy storage pool
- Copy active data to an active data pool
- Move data on a storage pool volume
- Migrate data from disk to sequential media
- Migrate data from sequential media to sequential media
- Back up, archive, or HSM migrate data that is initiated by client
- Reclaim volumes in a sequential-access storage pool
- Identify duplicates

Impacts of device changes on the SAN

The SAN environment can shift dramatically due to device or cabling changes. The dynamic nature of the SAN can cause static definitions to fail or become unpredictable.

Device IDs that are assigned by the SAN and known to the server or storage agent can be altered due to bus resets or other environmental changes. For example, the server might know a device X as *rmt0* (on AIX), based on the original path specification to the server and original configuration of the LAN. However, some event in the SAN, for example, the addition of new device Y, causes device X to be assigned *rmt1*. When the server tries to access device X by using *rmt0*, either the access fails or the wrong target device is accessed. The server attempts to recover from changes to devices on the SAN by using device serial numbers to confirm the identity of devices it contacts.

When you define a drive or library, you have the option of specifying the serial number for that device. If you do not specify the serial number when you define the device, the server obtains the serial number when you define the path for the device. In either case, the server then has the device serial number in its database and can use it to confirm the identity of a device for operations.

When the server uses drives and libraries on a SAN, the server attempts to verify that the correct device is used. The server contacts the device by using the device name in the path that you defined for it. The server then requests the serial number from the device, and compares that serial number with the serial number that is stored in the server database for that device.

If the serial number does not match, the server begins the process of discovery on the SAN, attempting to find the device with the matching serial number. If the server finds the device with the matching serial number, it corrects the definition of the path in the server's database by updating the device name in that path. The server issues a message with information about the change that is made to the device. Then, the server proceeds to use the device.

To determine when device changes on the SAN affect the IBM Storage Protect server, you can monitor the activity log for messages. The following messages are related to serial numbers:

- ANR8952 through ANR8958
- ANR8961 through ANR8968
- ANR8974 through ANR8975

Restriction: Some devices cannot report their serial numbers to applications such as the IBM Storage Protect server. If the server cannot obtain the serial number from a device, the server cannot help the system to recover from a device location change on the SAN.

Windows **Displaying device information**

You can display information about devices that are connected to the server by using the device information utility (tsmdlst).

Before you begin

- Ensure that the HBA API is installed. The HBA API is required to run the device information utility.
- Ensure that the tape device driver is installed and configured.

Procedure

1. From a command prompt, change to the server subdirectory of the server installation directory, for example, C:\Program Files\Tivoli\TSM\server.
2. Run the tsmdlst.exe executable file.

Related information

[QUERY SAN \(Query the devices on the SAN\)](#)
[tsmdlst \(Display information about devices\)](#)

Write-once, read-many tape media

Write-once, read-many (WORM) media help to prevent accidental or deliberate deletion of critical data. However, IBM Storage Protect imposes certain restrictions and guidelines to follow when you use WORM media.

You can use the following types of WORM media with IBM Storage Protect:

- IBM 3592, all supported generations
- IBM LTO-3 and all supported generations
- HP LTO-3 and all supported generations
- Quantum LTO-3 and all supported generations
- Quantum SDLT 600, Quantum DLT V4, and Quantum DLT S4
- StorageTek VolSafe
- Sony AIT50 and AIT100

Tips:

- A storage pool can consist of either WORM or RW media, but not both.
- To avoid wasting tape after a restore or import operation, do not use WORM tapes for database backup or export operations.

WORM-capable drives

To use WORM media in a library, all the drives in the library must be WORM-capable. A mount will fail if a WORM cartridge is mounted in a read/write (RW) drive.

However, a WORM-capable drive can be used as a RW drive if the WORM parameter in the device class is set to NO. Any type of library can have both WORM and RW media if *all* of the drives are WORM enabled. The only exception to this rule is NAS-attached libraries in which WORM tape media cannot be used.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Check-in of WORM media

The type of WORM media determines whether the media label needs to be read during check-in.

Library changers cannot identify the difference between standard read/write (RW) tape media and the following types of WORM tape media:

- VolSafe
- Sony AIT
- LTO
- SDLT
- DLT

To determine the type of WORM media that is being used, a volume must be loaded into a drive. Therefore, when you check in one of these types of WORM volumes, you must use the CHECKLABEL=YES option on the **CHECKIN LIBVOLUME** command.

If they provide support for WORM media, IBM 3592 library changers can detect whether a volume is WORM media without loading the volume into a drive. Specifying CHECKLABEL=YES is not required. Verify with your hardware vendors that your 3592 drives and libraries provide the required support.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Restrictions on WORM media

You cannot use prelabeled WORM media with the LTO or ECARTRIDGE device class.

You cannot use WORM media with IBM Storage Protect specified as the drive-encryption key manager for the following drives:

- IBM LTO-5, LTO-6, and later
- HP LTO-5, LTO-6, and later
- Oracle StorageTek T10000B
- Oracle StorageTek T10000C
- Oracle StorageTek T10000D

Mount failures with WORM media

If WORM tape media are loaded into a drive for a read-write (RW) device-class mount, it will cause a mount failure. Similarly, if RW tape media are loaded into a drive for a WORM device-class mount, the mount will fail.

Relabeling WORM media

You cannot relabel a WORM cartridge if it contains data. This applies to Sony AIT WORM, LTO WORM, SDLT WORM, DLT WORM, and IBM 3592 cartridges. The label on a VolSafe volume should be overwritten only once and only if the volume does not contain usable, deleted, or expired data.

Issue the **LABEL LIBVOLUME** command only once for VolSafe volumes. You can guard against overwriting the label by using the **OVERWRITE=NO** option on the **LABEL LIBVOLUME** command.

Related information

[LABEL LIBVOLUME \(Label a library volume\)](#)

Removing private WORM volumes from a library

If you perform an action on a WORM volume (for example, if you delete file spaces) and the server does not mark the volume as full, the volume is returned to scratch status. If a WORM volume is not marked as full and you delete it from a storage pool, the volume remains private. To remove a private WORM volume from a library, you must issue the **CHECKOUT LIBVOLUME** command.

Related information

[CHECKOUT LIBVOLUME \(Check a storage volume out of a library\)](#)

Creation of DLT WORM volumes

DLT WORM volumes can be converted from read/write (RW) volumes.

If you have SDLT-600, DLT-V4, or DLT-S4 drives and you want to enable them for WORM media, upgrade the drives by using V30 or later firmware available from Quantum. You can also use DLTIce software to convert unformatted RW volumes or blank volumes to WORM volumes.

In SCSI libraries, the IBM Storage Protect server creates scratch DLT WORM volumes automatically when the server cannot locate any scratch WORM volumes in a library's inventory. The server converts available unformatted or blank RW scratch volumes or empty RW private volumes to scratch WORM volumes. The server also rewrites labels on newly created WORM volumes by using the label information on the existing RW volumes.

Support for short and normal 3592 WORM tapes

IBM Storage Protect supports both short and normal 3592 WORM tapes. For best results, define them in separate storage pools

Querying a device class for the WORM-parameter setting

You can determine the setting of the WORM parameter for a device class by using the **QUERY DEVCLASS** command. The output contains a field, labeled WORM, and a value (YES or NO).

Related information

[QUERY DEVCLASS \(Display information on one or more device classes\)](#)


Troubleshooting problems with devices

You can troubleshoot errors that occur when you configure or use devices with IBM Storage Protect.

About this task

Use [Table 26 on page 132](#) to find a solution to the device-related problem.

Table 26. Resolving device problems

Symptom	Problem	Solution
Conflicts with other applications.	IBM Storage Protect requires a storage area network to share devices.	<p>Set up a storage area network.</p> <p> Attention: Data loss can occur if multiple IBM Storage Protect servers use the same device. Define or use a device with only one IBM Storage Protect server.</p> <p>Linux AIX Other applications can access IBM Storage Protect devices, by using a SCSI tape driver.</p>
Labeling fails.	A device for labeling volumes cannot be used at the same time that the server uses the device for other processes.	<p>You cannot overwrite existing volumes in a storage pool.</p> <p>You must resolve any hardware issues before you label a volume.</p>
	Incorrect or incomplete license registration.	Register the license for the device support that was purchased.
Conflicts among device drivers.	IBM Storage Protect issues messages about I/O errors when you define or use a sequential access device.	<p>Windows Windows device drivers and drivers that are provided by other applications can interfere with the IBM Storage Protect device driver if the IBM Storage Protect driver is not started first. To check on the order that device drivers are started by the system, complete the following steps:</p> <ol style="list-style-type: none"> 1. Click Control Panel. 2. Click Devices. Device drivers and their startup types are listed.
I/O errors	When you try to define or use a tape device, there might be device-driver conflicts. Windows device drivers and drivers that are provided by other applications can interfere with the IBM Storage Protect device driver if it is not started first.	
<p>Linux Unable to preempt tape drive reservation conflict with persistent reserve on Linux platform.</p>	<p>Linux On a Linux platform, the IBM Storage Protect server or storage agent requires that the IBM lin_tape device driver is configured for persistent reserve and an IBM pseudo device file /dev/TSMtape is created.</p>	<p>Linux If the Data Path failover is enabled in the IBM lin_tape driver, the /dev/TSMtape file is created automatically and persistent reserve can be used. Alternatively, configure persistent reserve for tape drive reservation on a Linux platform according to the following procedure:</p> <p>Tip: By default, the IBM lin_tape device driver uses SCSI-2 reserve to reserve tape drives.</p> <p>Linux</p> <ol style="list-style-type: none"> 1. Unload IBM lin_tape device driver. 2. In the lin_tape configuration file /etc/modprobe.conf or /etc/modprobe.conf.local (or, if you are running RHEL 6 or higher, the /etc/modprobe.d/lin_tape.conf), add the following line: <pre>options lin_tape tape_reserve_type=persistent</pre> 3. In the rules file /etc/udev/rules.d/98-lin_tape.rules, add the following line: <pre>KERNEL=="TSMtape", MODE="0666"</pre> 4. Reload IBM lin_tape device driver. <p>Linux The IBM pseudo file /dev/TSMtape is created and the IBM Storage Protect server can use persistent reserve to preempt tape drive reservation on Linux platforms.</p>

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 tape solution,” on page 135](#).

Part 3. Monitoring a tape solution

Monitor your tape-based solution to ensure correct operation.

About this task

After you implement your tape 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 email reports that summarize system status.

Procedure


1. Complete daily monitoring tasks. For instructions, see [Daily monitoring checklist](#).
2. Complete periodic monitoring tasks. For instructions, see [Periodic monitoring checklist](#).
3. Verify that your system complies with licensing requirements. For instructions, see [Verifying license compliance](#).
4. Optional: Set up email reports of system status. For instructions, see [“Tracking system status by using email reports”](#) on page 154

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 27. Daily monitoring tasks		
Task	Basic procedures	Advanced procedures and troubleshooting information
Watch for security notifications, which can indicate a ransomware attack.	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>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 27. 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 27. 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 27. 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: <div data-bbox="971 1087 1128 1113" data-label="Text"> <pre>query db f=d</pre> </div> 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 27. 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 27. 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 27. 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>Tape devices might have a warning or critical status if drives are unavailable. A drive is unavailable if it is offline, if it stopped responding to the server, or if its path is offline. A tape device might also have a critical status if the library is offline. Other columns of the Tape Devices table show the state of the library robotics, drives, and paths.</p> <p>To resolve issues with tape drives that have a critical state, you can take the drive offline if you need to use it for another activity, such as maintenance. To take a drive offline, complete the following steps:</p> <ol style="list-style-type: none"> 1. On the Operations Center Storage page, and select Tape Devices. 2. To view more information about a tape library, select a row and click Details. 3. To take a drive offline, select the tape drive and click Offline. <p>For tape backup operations, verify that sufficient scratch tapes are available. If you are not certain whether the number of available scratch tapes is sufficient, open the details notebook to view tape usage and an estimate of scratch tape availability. To open the details notebook, select a library in the table and click Details.</p>

Table 27. Daily monitoring tasks (continued)






Task	Basic procedures	Advanced procedures and troubleshooting information
<p>10 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 27. Daily monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting information
11 Monitor storage rules.	<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 operations run 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 28. 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. For information about this command, see 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>Limit the time for client backup operations to 8 - 12 hours. Ensure that client schedules do not overlap with server maintenance tasks.</p> <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 28. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Verify that current backup files for device configuration and volume history information are saved.</p>	<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: <div data-bbox="532 737 906 789" data-label="Text"> <pre>query option volhistory</pre> </div> <div data-bbox="532 804 906 856" data-label="Text"> <pre>query option devconfig</pre> </div> 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 28. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
<p>Determine whether sufficient space is available in the directory for the server instance.</p>	<p>Verify that at least 50 GB of free space is available in the directory for the server instance. 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> <p>where <i>instance_directory</i> specifies the instance directory.</p> 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> <p>where <i>instance_directory</i> specifies the instance directory.</p> 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"> Linux AIX /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 28. 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 28. 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 is 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 28. Periodic monitoring tasks (continued)

Task	Basic procedures	Advanced procedures and troubleshooting
Monitor and maintain tape devices.	<p>Monitor your environment for hardware errors on tape drives and tape libraries. For instructions, see “Monitoring tape alert messages for hardware errors” on page 151.</p> <p>Monitor media compatibility to prevent errors on tape drives. For instructions, see “Preventing errors caused by media incompatibility” on page 151.</p> <p>Monitor cleaning messages for tape drives. For instructions, see “Operations with cleaner cartridges” on page 152.</p>	
Evaluate the timing of client schedules. Ensure that the start and end times of client schedules do not overlap with server maintenance tasks. Limit the time for client backup operations to 8 - 12 hours.	<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 do not overlap with client schedules.	<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>The preferred method is to ensure that each maintenance task runs to completion before the next maintenance task starts. Examples of maintenance tasks include inventory expiration, copying of storage pools, space reclamation, and database backup.</p> <p>Tip: If a maintenance task is running too long, change the start time or the maximum runtime. 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 runtime for a task, issue the UPDATE SCHEDULE command. For information about this command, see UPDATE SCHEDULE (Update a client schedule).

Related information

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

Monitoring tape alert messages for hardware errors

Tape alert messages are generated by tape and library devices to report hardware errors. These messages help to determine problems that are not related to the server.

About this task

A log page is created and can be retrieved at any time or at a specific time such as when a drive is dismounted.

A tape alert message can have one of the following severity levels:

- Informational (for example, trying to load a cartridge type that is not supported)
- Warning (for example, a hardware failure is predicted)
- Critical (for example, there is a problem with the tape and data is at risk)

Tape alert messages are turned off by default.

Procedure

- To enable tape alert messages, issue the **SET TAPEALERTMSG** command and specify the **ON** value:
set tapealertmsg on
- To check whether tape alert messages are enabled, issue the **QUERY TAPEALERTMSG** command:
query tapealertmsg

Preventing errors caused by media incompatibility

By monitoring and resolving media compatibility issues, you can prevent errors in a tape-based solution. A new drive might have a limited ability to use media formats that are supported by a previous version of the drive. Often, a new drive can read but not write to the previous media format.

About this task

By default, existing volumes with a status of FILLING remain in that state after a drive upgrade. In some cases, you might want to continue to use a previous drive to fill these volumes. This preserves read/write capability for the existing volumes until they are reclaimed. If you choose to upgrade all of the drives in a library, verify that the media formats are supported by the new hardware. Unless you plan to use only the most current media with your new drive, you need to be aware of any compatibility issues. For migration instructions, see [“Migrating data to upgraded drives” on page 201](#).

To use a new drive with media that it can read but not write to, issue the **UPDATE VOLUME** command to set the access for those volumes to read-only. This prevents errors that are caused by read/write incompatibility. For example, a new drive might eject media that is written in a format that the drive does not support as soon as the media is loaded into the drive. Or a new drive might fail the first write command to media partially written in a format that the drive does not support.

When data on the read-only media expires and the volume is reclaimed, replace it with media that is fully compatible with the new drive. Errors can be generated if a new drive is unable to correctly calibrate a volume that is written when you use a previous format. To avoid this problem, ensure that the original drive is in good working order and at current microcode levels.

Operations with cleaner cartridges

To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Monitoring the cleaning process

If a cleaner cartridge is checked in to a library, and a drive must be cleaned, the server dismounts the data volume and runs the cleaning operation. If the cleaning operation fails or is canceled, or if no cleaner cartridge is available, you might not be aware that the drive needs cleaning. Monitor cleaning messages for these problems to ensure that drives are cleaned as needed. If necessary, issue the **CLEAN DRIVE** command to have the server try the cleaning again, or manually load a cleaner cartridge into the drive.

Using multiple cleaner cartridges

The server uses a cleaner cartridge for the number of cleanings that you specify when you check in the cleaner cartridge. If you check in two or more cleaner cartridges, the server uses only one of the cartridges until the designated number of cleanings for that cartridge is reached. Then, the server uses the next cleaner cartridge. If you check in two or more cleaner cartridges and issue two or more **CLEAN DRIVE** commands concurrently, the server uses multiple cartridges at the same time and decrements the remaining cleanings on each cartridge.

Related information

[AUDIT LIBRARY \(Audit volume inventories in an automated library\)](#)

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[CLEAN DRIVE \(Clean a drive\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

[QUERY LIBVOLUME \(Query a library volume\)](#)

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:</p> <p>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 capacity usage for replicated clients. For information about how to estimate capacity usage for these clients, see technote 1656476.</p> <p>a. On the Operations Center menu bar, hover over the settings icon  and click Licensing.</p> <p>b. Click the Back-end tab.</p> <p>c. Verify that the estimated amount of data complies with your licensing agreement.</p>
PVU model	<p>For information about how to assess compliance with PVU licensing terms, see Assessing compliance with the PVU licensing model.</p>

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.

What to do next

The general operations report includes an attachment. To find more detailed information, expand the sections in the attachment.

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

Part 4. Managing operations for a tape solution

Use this information to manage operations for a tape implementation for an IBM Storage Protect server.

Managing the Operations Center

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

About this task

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. For more information about managing the Operations Center, see [Managing the Operations Center](#).

Managing client operations

You can resolve client errors, manage client upgrades, and decommission client nodes that are no longer required. To free storage space on the server, you can deactivate obsolete data that is stored by application clients.

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](#).

For instructions about adding clients, see [“Protecting applications and systems” on page 107](#).

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

Optionally, 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](#).

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.

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

Software	Link to instructions
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, 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 \(Decommission a virtual machine\)](#)

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](#)

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” on page 13](#).
- 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 disk space for 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, see the information in [Resolving and preventing issues related to database growth and degraded performance in Tivoli Storage Manager V7.1.1.200 and later servers](#).

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.

- To increase or decrease the size of the active log, complete the following steps:
 - a) Ensure that the location for the active log has enough space for the increased log size.
 - b) Halt the server.
 - c) In the `dsmserv.opt` file, update the **ACTIVELOGSIZE** option to the new size of the active log, in megabytes.

The size of an active log file is based on the value of the **ACTIVELOGSIZE** option. Guidelines for space requirements are in the following table:

Table 29. How to estimate volume and file space requirements

ACTIVELOGSize option value	Reserve this much free space in the active log directory, in addition to the ACTIVELOGSize space
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 \(Set a server option for dynamic update\)](#)

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 backup and maintenance tasks are completing successfully. For more information about monitoring, see [Part 3, "Monitoring a tape solution,"](#) on page 135.
2. If the monitoring information shows that the server workload increased, you might need to 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 has performance issues. 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. Back up the database
 - b. Run expiration to remove client backups and archive file copies from server storage.

Related information

Deduplicating data (V7.1.1)

[Performance](#)

Optimizing operations by enabling collocation of client files

Collocation of client files reduces the number of volume mounts that are required when users restore, retrieve, or recall many files from a storage pool. Collocation thus reduces the amount of time that is required for these operations.

About this task

With collocation enabled, the server tries to keep files on a minimal number of sequential-access storage volumes. The files can belong to a single client node, a group of client nodes, a client file space, or a group of file spaces. You can set collocation for each sequential-access storage pool when you define or update the pool.

Figure 7 on page 164 shows an example of collocation by client node with three clients, each having a separate volume that contains that client's data.



Figure 7. Example of collocation enabled by node

Figure 8 on page 165 shows an example of collocation by group of client nodes. Three groups are defined, and the data for each group is stored on separate volumes.




Figure 8. Example of collocation enabled by node collocation group

Figure 9 on page 165 shows an example of collocation by file space group. Six groups are defined. Each group contains data from file spaces that belong to a single node. The data for each group is stored on a separate volume.




Figure 9. Example of collocation enabled by file space collocation group

When collocation is disabled, the server tries to use all available space on each volume before it selects a new volume. While this process provides better use of individual volumes, user files can become scattered across many volumes. Figure 10 on page 165 shows an example of collocation that is disabled, with three clients that share space on single volume.




Figure 10. Example of collocation disabled

With collocation disabled, more media mount operations might be required to mount volumes when users restore, retrieve, or recall many files.

Collocation by group is the IBM Storage Protect system default for primary sequential-access storage pools. The default for copy storage pools and retention storage pools is no collocation.

Effects of collocation on operations

The effect of collocation on resources and system performance depends on the type of operation that is being run.

Table 30 on page 166 summarizes the effects of collocation on operations.

Table 30. Effect of collocation on operations

Operation	Collocation enabled	Collocation disabled
Backing up, archiving, or migrating client files	More media mounts to collocate files.	Fewer media mounts are required.
Restoring, retrieving, or recalling client files	Large numbers of files can be restored, retrieved, or recalled more quickly because files are on fewer volumes.	Multiple mounts of media might be required for a single user because files might be spread across multiple volumes. More than one user's files can be stored on the same sequential-access storage volume. For example, if two users try to recover a file that is on the same volume, the second user is forced to wait until the first user's files are recovered.
Storing data on tape	The server attempts to use all available tape volumes to separate user files before it uses all available space on every tape volume.	The server attempts to use all available space on each tape volume before the server uses another tape volume.
Media mount operations	More mount operations are required when user files are backed up, archived, or migrated from client nodes directly to sequential-access volumes. More mount operations are required during reclamation and storage pool migration. More volumes are managed because volumes are not fully used.	More mount operations are required during restore, retrieve, and recall of client files.
Generating backup sets	Less time is spent searching database entries, and fewer mount operations are required.	More time is spent searching database entries and fewer mount operations are required.
Copying retention sets to tape Important: Your collocation setting can significantly increase the number of tape volumes that are needed by the retention set.	The server attempts to keep files from the same collocated entity on as few tape volumes as possible. The processing time to write a retention set to tape might increase.	The server attempts to use all available space on each tape volume before the server use another tape volume. If data needs to be restored from a retention set, more tape mounts might be required for a single retention set user because files might be spread across multiple volumes.

When collocation is enabled for a group, single client node, or file space, all the data that belongs to the group, the node, or the file space is moved or copied by one server process. For example, if data is collocated by group, all data for all nodes that belong to the same collocation group is migrated by the same process.

When collocating data, the IBM Storage Protect server tries to keep files together on a minimal number of sequential-access storage volumes. However, when the server is backing up data to volumes in a sequential-access storage pool, the backup process has priority over collocation settings. As a result, the server completes the backup operation, but might not be able to collocate the data.

For example, suppose that you are collocating by node and you specify that a node can use two mount points on the server. Suppose also that the data that is backed up from the node can easily fit on one tape volume. During backup, the server might mount two tape volumes, and the node's data might be distributed across two tapes, rather than one. If you enable collocation, the following server operations use one server process:

- Moving data from random-access and sequential-access volumes
- Moving node data from sequential-access volumes
- Backing up a random-access or sequential-access storage pool
- Restoring a sequential-access storage pool
- Reclaiming space in a sequential-access storage pool or offsite volumes
- Migrating data from a random-access storage pool

When you migrate data from a random-access disk storage pool to a sequential-access storage pool, and collocation is by node or file space, nodes or file spaces are automatically selected for migration based on the amount of data to be migrated. The node or file space with the most data is migrated first. If collocation is by group, all nodes in the storage pool are evaluated to determine which node has the most data. The node with the most data is migrated first along with all the data for all the nodes that belong to that collocation group. This process occurs, regardless of how much data is stored in the file spaces of nodes and regardless of whether the low migration threshold was reached.

However, when you migrate collocated data from a sequential-access storage pool to another sequential-access storage pool, the server orders the volumes according to the date when the volume was last accessed. The volume with the earliest access date is migrated first, and the volume with the latest access date is migrated last.

One reason to collocate by group is that individual client nodes often do not have sufficient data to fill high-capacity tape volumes. Collocating data by groups of nodes can reduce unused tape capacity by putting more collocated data on individual tapes. Also, collocating data by groups of file spaces reduces the unused tape to a greater degree.

The data that belongs to all the nodes in the same collocation group are migrated by the same process. Therefore, collocation by group can reduce the number of times that a volume to be migrated must be mounted. Collocation by group can also minimize database scanning and reduce tape passes during data transfer from one sequential-access storage pool to another.

Selecting volumes with collocation enabled

Volume selection depends on whether collocation is by group, node, or file space.

[Table 31 on page 168](#) shows how the IBM Storage Protect server selects the first volume when collocation is enabled for a storage pool at the client-node, collocation-group, and file-space level.

Table 31. How the server selects volumes when collocation is enabled

Volume Selection Order	When collocation is by group	When collocation is by node	When collocation is by file space
1	A volume that already contains files from the collocation group to which the client belongs	A volume that already contains files from the same client node	A volume that already contains files from the same file space of that client node
2	An empty predefined volume	An empty predefined volume	An empty predefined volume
3	An empty scratch volume	An empty scratch volume	An empty scratch volume
4	A volume with the most available free space among volumes that already contain data	A volume with the most available free space among volumes that already contain data	A volume that contains data from the same client node
5	Not applicable	Not applicable	A volume with the most available free space among volumes that already contain data

When the server must continue to store data on a second volume, it uses the following selection order to acquire more space:

1. An empty predefined volume
2. An empty scratch volume
3. A volume with the most available free space among volumes that already contain data
4. Any available volume in the storage pool

When collocation is by client node or file space, the server tries to provide the best use of individual volumes and minimizes file mixing from different clients or file spaces on volumes. This configuration is depicted in [Figure 11 on page 168](#), which shows that volume selection is *horizontal*, where all available volumes are used before all available space on each volume is used. A, B, C, and D represent files from four different client nodes.

Tips:

1. If collocation is by node and the node has multiple file spaces, the server does not attempt to collocate those file spaces.
2. If collocation is by file space and a node has multiple file spaces, the server attempts to put data for different file spaces on different volumes.




Figure 11. Using all available sequential-access storage volumes with collocation enabled at the node or file space level

Collocation can be by file space group or node group. When collocation is by node group (node collocation group), the server tries to collocate data from nodes that belong to the same collocation group. A file

space collocation group uses the same methods as a node collocation group, but can use more space because of the granularity of file space sizes. As shown in [Figure 12 on page 169](#), data for the following groups of nodes was collocated:

- Group 1 consists of nodes A, B, and C
- Group 2 consists of nodes D and E
- Group 3 consists of nodes F, G, H, and I

Whenever possible, the IBM Storage Protect server collocates data that belongs to a group of nodes on a single tape, as represented by Group 2 in the figure. Data for a single node can also be spread across several tapes that are associated with a group (Group 1 and 2). If the nodes in the collocation group have multiple file spaces, the server does not attempt to collocate those file spaces.




Figure 12. Using all available sequential-access storage volumes with collocation enabled at the group level

Normally, the IBM Storage Protect server always writes data to the current filling volume for the operation that is running. However, occasionally you might notice more than one filling volume in a collocated storage pool. Having more than one filling volume in a collocated storage pool can occur if different server processes or client sessions try to store data into the collocated pool at the same time. In this situation, IBM Storage Protect allocates a volume for each process or session that needs a volume so that both operations are completed as quickly as possible.

Selecting volumes with collocation disabled

When collocation is disabled, the server attempts to use all available space in a storage volume before it accesses another volume.

When you store client files in a sequential-access storage pool where collocation is disabled, the server selects a volume by using the following selection order:

1. A previously used sequential volume with available space (a volume with the most amount of data is selected first)
2. An empty volume

When the server needs to continue to store data on a second volume, it attempts to select an empty volume. If no empty volume exists, the server attempts to select any remaining available volume in the storage pool.

[Figure 13 on page 170](#) shows that volume use is vertical when collocation is disabled. In this example, fewer volumes are used because the server attempts to use all available space by mixing client files on individual volumes. A, B, C, and D represent files from four different client nodes.




Figure 13. Using all available space on sequential-access volumes with collocation disabled

Collocation settings

After you define a storage pool, you can change the collocation setting by updating the storage pool. The change in collocation for the pool does not affect files that are already stored in the pool.

For example, if collocation is off for a storage pool and you turn it on, from then on client files that are stored in the pool are collocated. Files that were previously stored in the storage pool are not moved to collocate them. As volumes are reclaimed or restored, the data in the pool tends to become more collocated. You can also use the **MOVE DATA** or **MOVE NODEDATA** commands to move data to new volumes to increase collocation. Moving data to new volumes causes an increase in the processing time and the volume mount activity.

Tip: A mount wait can occur or take longer than usual when collocation by file space is enabled and a node has a volume that contains multiple file spaces. If a volume is eligible to receive data, IBM Storage Protect waits for that volume.

Collocation of copy storage pools

Using collocation on copy storage pools requires special consideration. Collocation of copy storage pools, especially by node or file space, results in more partially filled volumes and potentially unnecessary offsite reclamation activity.

Primary storage pools play a different recovery role than copy storage pools. Normally, you use primary storage pools to recover data to clients directly. In a disaster, when both clients and the server are lost, you might use offsite copy storage pool volumes to recover the primary storage pools. The types of recovery scenarios can help you to determine whether to use collocation on your copy storage pools.

Collocation typically results in partially filled volumes when you collocate by node or by file space. However, partially filled volumes are less prevalent when you collocate by group. Partially filled volumes might be acceptable for primary storage pools because the volumes remain available and can be filled during the next migration process. However, partially filled volumes might be unacceptable for copy storage pools whose storage pool volumes are taken offsite immediately. If you use collocation for copy storage pools, you must make the following decisions:

- Taking more partially filled volumes offsite, which increases the reclamation activity when the reclamation threshold is lowered or reached.
- Leaving these partially filled volumes onsite until they fill and risk not having an offsite copy of the data on these volumes.
- Whether to collocate by group to use as much tape capacity as possible.

When collocation is disabled for a copy storage pool, typically only a few partially filled volumes remain after data is backed up to the copy storage pool.

Consider your options carefully before you use collocation for copy storage pools, and whether to use simultaneous write. If you do not use simultaneous write and you use collocation for your primary storage pools, you might want to disable collocation for copy storage pools. Collocation of copy storage pools might be desirable if you have few clients with each of them having large amounts of incremental backup

data each day. For collocation with simultaneous write, you must ensure that the collocation settings are identical for the primary storage pools and copy storage pools.

Collocation of retention storage pools

The value that you select for the collocation property affects how a retention set's data is distributed across tape volumes. In general, to minimize the number of tape volumes, disable collocation. By default, the collocation setting for retention storage pools is disabled.

When you disable the collocation setting, the server attempts to use all available space on each tape volume before it selects a new volume. Although this process makes more efficient use of individual tape volumes, the data for each retention set is not collocated and might be distributed across many tape volumes.

Performance considerations: Your collocation settings can have a significant impact on system performance. System performance might decrease when the retention set data is written to tape and during restore operations. Before you enable collocation settings for retention storage pools, consider your organization's requirements and the performance tradeoffs. If collocation is enabled, the server tries to keep files for each entity on a minimal number of tape volumes. However, this option increases both the server processing time to collocate the files and the number of required tape volumes.

Where the system copies data is also affected by the **STACK** parameter value that you specify for a retention set:

Yes

The retention set data can share tape volumes with data that is copied from other retention sets. During volume selection, the first volumes to be filled are those that are already in a FILLING state, but only if those volumes are not already in use by retention sets that require a separate volume.

No

The retention set is collocated on one or more tape volumes and data from other retention sets is not placed on those volumes. Empty volumes are filled first, but data can also be copied to FILLING volumes if the volumes already contain data for the retention set that is being copied.

Restrictions:

- The **STACK** parameter applies only when you copy retention data to tape volumes. The parameter is ignored when you copy retention data to cloud storage.
- The collocation setting affects only the data that is written to the retention storage pool after the setting is changed. Files that are already stored in the pool are not affected.

You can enable collocation or change collocation settings by specifying the **COLLOCATE** parameter on the **DEFINE STGPOOL** or **UPDATE STGPOOL** commands. Empty volumes are filled first,

Related concepts

[“Selecting volumes with collocation disabled” on page 169](#)

When collocation is disabled, the server attempts to use all available space in a storage volume before it accesses another volume.

[“Effects of collocation on operations” on page 166](#)

The effect of collocation on resources and system performance depends on the type of operation that is being run.

Planning for and enabling collocation

Understanding the effects of collocation can help reduce the number of media mounts, make better use of space on sequential volumes, and improve the efficiency of server operations.

About this task

Table 32 on page 172 lists the four collocation options that you can specify on the **DEFINE STGPOOL** and **UPDATE STGPOOL** commands. The table also shows the effects of collocation on data that belongs to nodes that are and are not members of collocation groups.

Table 32. Collocation options and the effects on node data

Collocation option	If a node is not defined as a member of a collocation group	If a node is defined as a member of a collocation group
No	The data for the node is not collocated.	The data for the node is not collocated.
Group	The server stores the data for the node on as few volumes in the storage pool as possible.	The server stores the data for the node and for other nodes that belong to the same collocation group on as few volumes as possible.
Node	The server stores the data for the node on as few volumes as possible.	The server stores the data for the node on as few volumes as possible.
File space	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.

Table 33. Collocation group options and effects on file space data

Collocation option	If a file space is not defined as a member of a collocation group	If a file space is defined as a member of a collocation group
No	The data for the file space is not collocated.	The data for the file space is not collocated.
Group	The server stores the data for the file space on as few volumes in the storage pool as possible.	The server stores the data for the file space and other file spaces that belong to the same collocation group on as few volumes as possible.
Node	The server stores the data for the node on as few volumes as possible.	The server stores the data for the node on as few volumes as possible.
File space	The server stores the data for the node's file space on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.	The server stores the data for the file spaces on as few volumes as possible. If a node has multiple file spaces, the server stores the data for different file spaces on different volumes in the storage pool.

Procedure

To determine whether and how to collocate data, complete the following steps:

1. Determine how to organize data, whether by client node, group of client nodes, or file space. To collocate by group, you must decide how to group nodes:
 - If the goal is to save space, you might want to group small nodes together to better use tapes.
 - If the goal is potentially faster client restores, group nodes together so that they fill as many tapes as possible. By grouping nodes together, the individual node data is distributed across two or more tapes and that more tapes can be mounted simultaneously during a multi-session no-query restore operation.
 - If the goal is to departmentalize data, you can group nodes by department.
2. To collocate groups, complete the following steps:
 - a) Define collocation groups with the **DEFINE COLLOGROUP** command.
 - b) Add client nodes to the collocation groups with the **DEFINE COLLOCMEMBER** command.

The following query commands are available to help in collocating groups:

QUERY COLLOCGROUP

Displays the collocation groups that are defined on the server.

QUERY NODE

Displays the collocation group, if any, to which a node belongs.

QUERY NODEDATA

Displays information about the data for one or more nodes in a sequential-access storage pool.

QUERY STGPOOL

Displays information about the location of client data in a sequential-access storage pool and the amount of space a node occupies in a volume.

You can also use IBM Storage Protect server scripts or PerlL scripts to display information that can be useful in defining collocation groups.

3. Specify how data must be collocated in a storage pool by issuing the **DEFINE STGPOOL** or **UPDATE STGPOOL** command and specifying the **COLLOCATE** parameter.

What to do next

Tip: To reduce the number of media mounts, use space on sequential volumes more efficiently, and enable collocation, complete the following steps:

- Define a storage pool hierarchy and policy to require that backed-up, archived, or space-managed files are initially stored in disk storage pools.

When files are migrated from a disk storage pool, the server attempts to migrate all files that belong to the client node or collocation group that is using the most disk space in the storage pool. This process works well with the collocation option because the server tries to place all of the files from a particular client on the same sequential-access storage volume.

- Use scratch volumes for sequential-access storage pools to allow the server to select new volumes for collocation.
- Specify the client option COLLOCATEBYFILESPEC to limit the number of tapes to which objects associated with one file specification are written. This collocation option makes collocation by the server more efficient; it does not override collocation by file space or collocation by node.

Managing tape devices

Routine tape operations include preparing tape volumes for use, controlling how and when volumes are reused, and ensuring that sufficient volumes are available. You also must respond to operator requests and manage libraries, drives, disks, paths, and data movers.

Preparing removable media

You must prepare removable media before it can be used to store data. Typical preparation tasks include labeling and checking in volumes.

About this task

When IBM Storage Protect accesses a removable media volume, it verifies the volume name in the label header to ensure that the correct volume is accessed.

Tape volumes must be labeled before the server can use them.

Procedure

To prepare a volume for use, complete the following steps:

1. Label the volume by issuing the **LABEL LIBVOLUME** command.

2. For automated libraries, check the volume into the library. For instructions, see [“Checking volumes into an automated library”](#) on page 176,

Tip: When you use the **LABEL LIBVOLUME** command with drives in an automated library, you can label and check in the volumes with one command.

3. If the storage pool cannot contain scratch volumes (**MAXSCRATCH=0**), identify the volume to IBM Storage Protect by name so that the volume can be accessed later.

If the storage pool can contain scratch volumes (**MAXSCRATCH** is set to a non-zero value), skip this step.

Labeling tape volumes

You must label tape volumes before the server can use them.

About this task

For automated libraries, you are prompted to insert the volume in the entry/exit slot of the library. If no convenience input/output (I/O) station is available, insert the volume into an empty slot. You can label the volumes when you check them in or before you check them in.

Procedure

To label tape volumes before you check them in, complete the following steps:

1. Label tape volumes by issuing the **LABEL LIBVOLUME** command.
For example, to name a library volume **VOLUME1** in a library that is named **LIBRARY 1**, issue the following command:

```
label libvolume library1 volume1
```

Requirement: At least one drive must be available. The drive cannot be used by another IBM Storage Protect process. If a drive is idle, the drive is considered to be unavailable.

2. To overwrite an existing label, specify the **OVERWRITE=YES** parameter. By default, the **LABEL LIBVOLUME** command does not overwrite an existing label.

Related tasks

[Labeling new volumes by using AUTOLABEL](#)

Using the **AUTOLABEL** parameter on the **DEFINE LIBRARY** or **UPDATE LIBRARY** command is more efficient than using the **LABEL LIBVOLUME** command, which requires you to mount volumes separately.

Related information

[LABEL LIBVOLUME \(Label a library volume\)](#)

Labeling volumes in a SCSI library

You can label volumes individually or use IBM Storage Protect to search the library for volumes and label the found volumes.

Labeling volumes individually

When you label volumes individually by using the **LABEL LIBVOLUME** command, you must specify a volume name.

Procedure

1. Insert volumes into the entry/exit slot of the library when the server prompts you. The library mounts each inserted volume into a drive.
2. For a SCSI library, enter a volume name when prompted. A label with the specified name is written to the volume.

Tip: To prompt for the volume name for a SCSI library, issue the **LABEL LIBVOLUME** command and specify the **LABELSOURCE=PROMPT** parameter.

3. If the library does not have an entry/exit port, you are prompted to remove the tape from a specified slot number. Remove the tape from the specified slot.

If the library has an entry/exit port, the command by default returns each labeled volume to the entry/exit port of the library.

Overwriting volume labels in a SCSI library

You can use the **LABEL LIBVOLUME** command to overwrite existing volume labels if no valid data exists in the storage volumes.

About this task

You can label volumes in a SCSI library, even if they do not have an entry/exit port. You must manually insert each new volume to the library, and place the volumes in storage slots inside the library after their labels are written.

Procedure

Overwrite the existing volume labels by issuing the **LABEL LIBVOLUME** command. For example, if the name of the library is LIB1 and the volume name is VOLNAME, issue the following command:

```
label libvolume lib1 volname overwrite=yes
```

Labeling new volumes by using AUTOLABEL

Using the **AUTOLABEL** parameter on the **DEFINE LIBRARY** or **UPDATE LIBRARY** command is more efficient than using the **LABEL LIBVOLUME** command, which requires you to mount volumes separately.

Procedure

Issue the **DEFINE LIBRARY** or **UPDATE LIBRARY** command and specify the **AUTOLABEL** parameter.

Tip: If you use the **AUTOLABEL** parameter with a SCSI library, you must check in tapes by specifying **CHECKLABEL=BARCODE** parameter on the **CHECKIN LIBVOLUME** command. The **AUTOLABEL** parameter defaults to YES for all non-SCSI libraries and to NO for SCSI libraries. The **CHECKLABEL=BARCODE** parameter is honored only if the library has a bar code reader.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

Searching a library and labeling volumes

IBM Storage Protect can search all storage slots in a library for volumes and can attempt to label each volume that it finds.

Procedure

To search a library and label volumes, issue the **LABEL LIBVOLUME** command and specify the **SEARCH=YES** parameter.

Tip: If you use a SCSI library and the library has a bar code reader, the **LABEL LIBVOLUME** command can use the reader to obtain volume names, instead of prompting you for volume names. The **LABELSOURCE=BARCODE** parameter is valid only for SCSI libraries.

For example, to label all volumes in a SCSI library, issue the following command:

```
label libvolume library_name search=yes labelsource=barcode
```

IBM Storage Protect selects the next available drive so that you can continue your search.

Results

After a volume is labeled, the volume is returned to its original location in the library.

Related information

[LABEL LIBVOLUME \(Label a library volume\)](#)

Checking volumes into an automated library

You can check in a volume to an automated library by using the **CHECKIN LIBVOLUME** command.

Before you begin

To automatically label tapes before you check them in, issue the **DEFINE LIBRARY** command and specify the **AUTOLABEL=YES** parameter. By using the **AUTOLABEL** parameter, you eliminate the need to prelabel a set of tapes.

About this task

Each volume that is used by a server for any purpose must have a unique name. This requirement applies to all volumes, whether the volumes are used for storage pools, or used for operations such as database backup or export. The requirement also applies to volumes that are in different libraries but that are used by the same server.

Tips:

- Do not use a single library for volumes that have bar code labels and volumes that do not have bar code labels. Bar code scanning can take a long time for unlabeled volumes.
- The server accepts only tapes that are labeled with IBM standard labels.
- Any volume that has a bar code that begins with CLN is treated as a cleaning tape.
- If a volume has an entry in volume history, you cannot check it in as a scratch volume.

Procedure

1. To check a storage volume into a library, issue the **CHECKIN LIBVOLUME** command.

Tip: The command always runs as a background process. Wait for the **CHECKIN LIBVOLUME** process to complete processing before you define volumes, or the defining process fails. You can save time by checking in volumes as part of the labeling operation.

2. Name the library and specify whether the volume is a private volume or a scratch volume. Depending on whether you use scratch volumes or private volumes, complete one of the following steps:
 - If you use only scratch volumes, ensure that enough scratch volumes are available. For example, you might need to label more volumes. As volumes are used, you might also need to increase the number of scratch volumes that are allowed in the storage pool that you defined for this library.
 - If you want to use private volumes in addition to or instead of scratch volumes in the library, define volumes to the storage pool by using the **DEFINE VOLUME** command. You must label and check in the volumes that you define.

Related tasks

[Labeling tape volumes](#)

You must label tape volumes before the server can use them.

Checking a single volume into a SCSI library

You can check in a single volume by issuing the **CHECKIN LIBVOLUME** command and specifying the **SEARCH=NO** parameter. IBM Storage Protect requests that the mount operator load the volume into the entry/exit port of the library.

Procedure

1. Issue the **CHECKIN LIBVOLUME** command.

For example, to check in volume VOL001, enter the following command:

```
checkin libvolume tapelib vol001 search=no status=scratch
```

2. Respond to the prompt from the server.

- If the library has an entry/exit port, you are prompted to insert a tape into the entry/exit port.
- If the library does not have an entry/exit port, you are prompted to insert a tape into one of the slots in the library. Element addresses identify these slots. For example, the server finds that the first empty slot is at element address 5. The following message is returned:

```
ANR8306I 001: Insert 8MM volume VOL001 R/W in slot with element  
address 5 of library TAPELIB within 60 minutes; issue 'REPLY' along  
with the request ID when ready.
```

If you do not know the location of element address 5 in the library, check the worksheet for the device. To find the worksheet, review the documentation for your library. After you insert the volume as requested, respond to the message from an IBM Storage Protect administrative client. Issue the **REPLY** command, followed by the request number (the number at the beginning of the mount request) for example:

```
reply 1
```

Tip: Element addresses are sometimes numbered starting with a number other than 1. Check the worksheet to be sure. If no worksheet is listed for your device in [IBM Support Portal for IBM Storage Protect](#), see the documentation for your library.

If you specify a wait time of 0 by using the optional **WAITTIME** parameter on the **CHECKIN LIBVOLUME** command, a **REPLY** command is not required. The default wait time is 60 minutes.

Checking in volumes from library storage slots

When you have many volumes to check in and you want to avoid issuing a **CHECKIN LIBVOLUME** command for each volume, you can search storage slots for new volumes. The server finds volumes that have not yet been added to the volume inventory.

Procedure

1. Open the library and place the new volumes in unused slots.
For example, for a SCSI device, open the library access door, place all of the new volumes in unused slots, and close the door.
2. If the volumes are not labeled, use the **LABEL LIBVOLUME** command to label the volume.
3. Issue the **CHECKIN LIBVOLUME** command with the **SEARCH=YES** parameter.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Checking in volumes from library entry/exit ports

You can search all slots of bulk entry/exit ports for labeled volumes and the server can check them in automatically.

Before you begin

Issue the **LABEL LIBVOLUME** command to label volumes that are not labeled.

About this task

For SCSI libraries, the server scans all of the entry/exit ports in the library for volumes. If a volume is found that contains a valid volume label, it is checked in automatically.

Procedure

Issue the **CHECKIN LIBVOLUME** command and specify the **SEARCH=BULK** parameter.

- To load a tape in a drive and read the label, specify the **CHECKLABEL=YES** parameter. After the server reads the label, the server moves the tape from the drive to a storage slot.
- To have the server use the bar code reader to verify external labels on tapes, specify the **CHECKLABEL=BARCODE** parameter. When bar code reading is enabled, the server reads the label and moves the tape from the entry/exit port to a storage slot.

Checking in volumes by using library bar code readers

You can save time when you check in volumes to libraries that have bar code readers by using the characters on the bar code labels as names for the volumes.

About this task

The server reads the bar code labels and uses the information to write the internal media labels. For volumes that have no bar code labels, the server mounts the volumes in a drive and attempts to read the internal, recorded label.

Procedure

Issue the **CHECKIN LIBVOLUME** command with the **CHECKLABEL=BARCODE** parameter.

For example, to use a bar code reader to search a library that is named TAPELIB and check in a scratch tape, issue the following command:

```
checkin libvolume tapelib search=yes status=scratch checklabel=barcode
```

Checking in volumes by using a bar code reader

You can save time when you check in volumes by using a bar code reader, if your library has one.

About this task

When you check in a volume, you can specify whether the media labels are read during check-in processing. When label-checking is on, IBM Storage Protect mounts each volume to read the internal label and checks in a volume only if it is correctly labeled. Label-checking can prevent future errors when volumes are used in storage pools, but also increases processing time at check-in.

If a volume has no bar code label, IBM Storage Protect mounts the volumes in a drive and attempts to read the recorded label.

Procedure

To check in volumes by using a bar code reader, issue the **CHECKIN LIBVOLUME** command and specify **CHECKLABEL=BARCODE**. For example, to use the bar code reader to check in all volumes as scratch volumes in a library that is named TAPELIB, issue the following command:

```
checkin libvolume tapelib search=yes status=scratch checklabel=barcode
```

Related tasks

[Preparing removable media](#)

You must prepare removable media before it can be used to store data. Typical preparation tasks include labeling and checking in volumes.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Checking volumes into a full library with swapping

If no empty slots are available in the library when you are checking in volumes, the check-in operation fails unless you enable *swapping*. If you enable swapping and the library is full, the server selects a volume to eject and then checks in the volume that you requested.

About this task

The server selects the volume to eject by checking first for any available scratch volume, then for the volume that is least frequently mounted. The server ejects the volume that it selects for the swap operation from the library and replaces the ejected volume with the volume that is being checked in.

Procedure

- To swap volumes if an empty library slot is not available to check in a volume, issue the **CHECKIN LIBVOLUME** command and specify the **SWAP=YES** parameter.
For example, to check in a volume that is named VOL1 into a library that is named AUTO and specify swapping, issue the following command:

```
checkin libvolume auto vol1 swap=yes
```

Related tasks

[Managing a full library with an overflow location](#)

As the demand for storage grows, the number of volumes that you need for a storage pool might exceed the physical capacity of an automated library. To make space available for new volumes and to monitor existing volumes, you can define an overflow location for a storage pool.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Private volumes and scratch volumes

To optimize tape storage, review the information about private volumes and scratch volumes. Use private volumes and scratch volumes appropriately.

Private volumes cannot be overwritten when a scratch mount is requested. You cannot check in a volume with scratch status when that volume is used by a storage pool, to export data, to back up a database or to back up to a backup set volume.

Partially written volumes are always private volumes. Volumes have a status of either scratch or private, but when IBM Storage Protect stores data on them, their status becomes private.

Table 34. Private volume and scratch volume uses

Type of volume	When to use
Private volumes	Use private volumes to regulate the volumes that are used by individual storage pools, and to manually control the volumes. To define private volumes, issue the DEFINE VOLUME command. For database restore, memory dumps, or loads, or for server import operations, you must specify private volumes.
Scratch volumes	<p>In some cases, you can simplify volume management by using scratch volumes. You can use scratch volumes in the following circumstances:</p> <ul style="list-style-type: none"> • When you do not need to define each storage pool volume. • When you want to take advantage of the automation of robotic devices. • When different storage pools share an automated library, and the storage pools can dynamically acquire volumes from the scratch volumes in the library. The volumes do not have to be preallocated to the storage pools.

Related tasks

[Changing the status of a volume in an automated library](#)

You can change the status of a volume from private to scratch or from scratch to private.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[DELETE VOLUME \(Delete a storage pool volume\)](#)

Element addresses for library storage slots

An element address is a number that indicates the physical location of a storage slot or drive within an automated library.

If a library has entry/exit ports, you can add and remove media by using the ports. If no entry/exit port exists, you must load tapes into storage slots.

If you load tapes into storage slots, you must reply to mount requests that identify storage slots with element addresses. If you specify a wait time of 0 on the **CHECKIN LIBVOLUME** command or the **LABEL LIBVOLUME** command, you do not need to reply to a mount request.

For element addresses, see the device manufacturer's documentation or go to the [IBM Support Portal for IBM Storage Protect](#) and search for element addresses.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

Managing volume inventory

You can manage volume inventory by controlling the server's access to volumes, by reusing tapes, and by reusing volumes that are used for database backup and export operations. You can also manage inventory by maintaining a supply of scratch volumes.

About this task

Each volume that is used by a server must have a unique name, whether the volumes are used for storage pools, or used for operations such as database backup or export. Volumes that are in different libraries but that are used by the same server must also have a unique name.

Controlling access to volumes

You can use different methods to control access to volumes.

Procedure

To control access to volumes, take any of the following actions:

- To prevent the server from mounting a volume, issue the **UPDATE VOLUME** command and specify the **ACCESS=UNAVAILABLE** parameter.
- To make volumes unavailable and send them offsite for protection, use a copy storage pool or an active-data storage pool.
- You can back up primary storage pools to a copy storage pool and then send the copy storage pool volumes offsite.
- You can copy active versions of client backup data to active-data storage pools, and then send the volumes offsite.
- You can track copy storage pool volumes and active-data pool volumes by changing their access mode to offsite, and updating the volume history to identify their location.

Related information

[UPDATE VOLUME \(Update a storage pool volume\)](#)

Reusing tapes

To ensure an adequate supply of tapes, you can expire old files, reclaim volumes, and delete volumes that reach end of life. You can also maintain a supply of scratch volumes.

About this task

Over time, media age, and you might not need some of the backup data that is stored on the media. You can define server policies to determine how many backup versions are retained and how long they are retained. You can use expiration processing to delete files that you no longer require. You can keep the data that you require on the media. When you no longer require the data, you can then reclaim and reuse the media.

Procedure

1. Delete unnecessary client data by regularly running expiration processing. Expiration processing deletes data that is no longer valid either because it exceeds the retention specifications in the policy or because users or administrators deleted the active versions of the data.
2. Reuse volumes in storage pools by running reclamation processing.

Reclamation processing consolidates any unexpired data by moving it from multiple volumes onto fewer volumes. The media can then be returned to the storage pool and reused.
3. Reuse volumes that contain outdated database backups or exported data that is no longer required by deleting volume history.

Before the server can reuse volumes that are tracked in the volume history, you must delete the volume information from the volume history file by issuing the **DELETE VOLHISTORY** command.

Tip: If your server uses the disaster recovery manager (DRM) function, the volume information is automatically deleted during **MOVE DRMEDIA** command processing.

4. Determine when tape volumes reach end of life. You can use the server to display statistics about volumes, including the number of write operations that are completed on the media and the number of write errors. Private volumes and scratch volumes display the following statistical data:

Private volumes

For media initially defined as private volumes, the server maintains this statistical data, even as the volume is reclaimed. You can compare the information with the number of write operations and write errors that are recommended by the manufacturer.

Scratch volumes

For media initially defined as scratch volumes, the server overwrites this statistical data each time the volumes are reclaimed.

5. Reclaim any valid data from volumes that reach end of life. If the volumes are in automated libraries, check them out of the volume inventory. Delete private volumes from the database with the **DELETE VOLUME** command.
6. Ensure that volumes are available for tape rotation so that the storage pool does not run out of space. You can use the Operations Center to monitor the availability of scratch volumes. Ensure that the number of scratch volumes is high enough to meet demand. For more information, see [“Maintaining a supply of volumes in a library that contains WORM media”](#) on page 183.

WORM media: Write Once Read Many (WORM) drives can waste media when the server cancels transactions because volumes are unavailable to complete the backup operation. After the server writes to WORM volumes, the space on the volumes cannot be reused, even if the transactions are canceled (for example, if a backup is canceled because of a shortage of media in the device). To minimize wasted WORM media, complete the following actions:

- a. Ensure that the maximum number of scratch volumes for the device storage pool is at least equal to the number of storage slots in the library.
- b. Check enough volumes into the device's volume inventory for the expected load.

If most backup operations are for small files, controlling the transaction size can affect how WORM platters are used. Smaller transactions mean that less space is wasted when a transaction such as a backup operation must be canceled. Transaction size is controlled by a server option, **TXNGROUPMAX**, and a client option, **TXNBYTELIMIT**.

Related tasks

Migrating data to upgraded drives

If you upgrade all of the tape drives in a library, you can preserve your existing policy definitions to migrate and expire existing data, and you can use the new drives to store data.

Managing server requests for volumes

IBM Storage Protect displays requests and status messages to all administrative command-line clients that are started in console mode. These request messages often have a time limit. Successful server operations must be completed within the time limit that is specified; otherwise, the operation times out.

Related information

[DELETE VOLHISTORY \(Delete sequential volume history information\)](#)

[DELETE VOLUME \(Delete a storage pool volume\)](#)

[EXPIRE INVENTORY \(Manually start inventory expiration processing\)](#)

[RECLAIM STGPOOL \(Reclaim volumes in a sequential-access storage pool\)](#)

[Txnbytelimit option](#)

[TXNGROUPMAX server option](#)

Maintaining a supply of scratch volumes

You must set the maximum number of scratch volumes for a storage pool high enough for the expected usage.

About this task

When you define a storage pool, you must specify the maximum number of scratch volumes that the storage pool can use. The server automatically requests a scratch volume when needed. When the number of scratch volumes that the server is using for the storage pool exceeds the specified maximum, the storage pool can run out of space.

Procedure

When a storage pool needs more than the maximum number of scratch volumes, you can take one or both of the following actions:

1. Increase the maximum number of scratch volumes by issuing the **UPDATE STGPOOL** command and specifying the **MAXSCRATCH** parameter.
2. Make volumes available for reuse by running expiration processing and reclamation to consolidate data onto fewer volumes.

- a) Issue the **EXPIRE INVENTORY** command to run expiration processing.

Tip: By default this process automatically runs every day. You can also specify the **EXPINTERVAL** server option in the server options file, `dsmserv.opt`, to run expiration processing automatically. A value of 0 means that you must use the **EXPIRE INVENTORY** command to run expiration processing.

- b) Issue the **RECLAIM STGPOOL** command to run reclamation processing.

Tip: You can also specify reclamation thresholds when you define the storage pool by using the **DEFINE STGPOOL** command and specifying the **RECLAIMPROCESS** parameter.

What to do next

If you need more volumes for future backup operations, label more scratch volumes by using the **LABEL LIBVOLUME** command.

Related tasks

[Maintaining a supply of scratch volumes in an automated library](#)

When you define a storage pool that is associated with an automated library, you can specify a maximum number of scratch volumes equal to the physical capacity of the library. If the server is using a greater number of scratch volumes for the storage pool, you must ensure that enough volumes are available.

Related information

[EXPIRE INVENTORY \(Manually start inventory expiration processing\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

[RECLAIM STGPOOL \(Reclaim volumes in a sequential-access storage pool\)](#)

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

Maintaining a supply of volumes in a library that contains WORM media

For libraries that contain Write Once Read Many (WORM) media, you can prevent cancellation of data storage transactions by maintaining a supply of scratch or new private volumes in the library. Canceled transactions can cause WORM media to be wasted.

About this task

IBM Storage Protect cancels a transaction if volumes, either private or scratch, are unavailable to complete the data storage operation. After IBM Storage Protect begins a transaction by writing to a WORM volume, the written space on the volume cannot be reused, even if the transaction is canceled.

For example, if you have WORM volumes that hold 2.6 GB each and a client starts to back up a 12 GB file. If IBM Storage Protect cannot acquire a fifth scratch volume after four volumes are full, IBM Storage Protect cancels the backup operation. The four volumes that IBM Storage Protect already filled cannot be reused.

To minimize cancellation of transactions, you must have enough volumes available in the library to manage expected client operations such as backups.

Procedure

1. Ensure that the storage pool that is associated with the library has sufficient scratch volumes. Issue the **UPDATE STGPOOL** command and specify the **MAXSCRATCH** parameter.
2. To manage the expected load, check in a sufficient number of scratch or private volumes to the library by issuing the **CHECKIN LIBVOLUME** command.
3. To control transaction size, specify the TXNGROUPMAX server option and the TXNBYTELIMIT client option. If your clients tend to store small files, controlling the transaction size can affect how WORM volumes are used. Smaller transactions waste less space when a transaction such as a backup must be canceled.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

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

[Txnbytelimit option](#)

[TXNGROUPMAX server option](#)

Manage the volume inventory in automated libraries

The IBM Storage Protect server uses a library volume inventory to track scratch and private volumes that are available in an automated library. You must ensure that the inventory is consistent with the volumes that are physically in the library.

The library volume inventory is separate from the inventory of volumes for each storage pool. To add a volume to a library volume inventory, you check in a volume to that IBM Storage Protect library.

A list of volumes in the library volume inventory might not be identical to a list of volumes in the storage pool inventory for the device. For example, you can check in scratch volumes to the library but you cannot define them to a storage pool. If scratch volumes are not selected for backup operations, you can define private volumes to a storage pool but you cannot check them into the volume inventory for the device.

To ensure that the volume inventory for the server library remains accurate, check out volumes to physically remove the volumes from a SCSI library. When you check out a volume that is used by a storage pool, the volume remains in the storage pool. If you must mount the volume when it is checked out, a message to the mount operator's console is displayed with a request to check in the volume. If the check-in operation is unsuccessful, the server marks the volume as unavailable.

When a volume is in the library volume inventory, you can change the status of the volume from scratch to private.

To check whether the volume inventory for the server library is consistent with the volumes that are physically in the library, you can audit the library. The inventory can become inaccurate if volumes are moved in and out of the library without informing the server by using volume check-in or check-out operations.

Related tasks

[Checking volumes into an automated library](#)

You can check in a volume to an automated library by using the **CHECKIN LIBVOLUME** command.

Related information

[AUDIT LIBRARY \(Audit volume inventories in an automated library\)](#)

Changing the status of a volume in an automated library

You can change the status of a volume from private to scratch or from scratch to private.

Procedure

To change the status of a volume, issue the **UPDATE LIBVOLUME** command.

For example, to change the status of a volume that is named VOL1 to a private volume, issue the following command:

```
update libvolume lib1 vol1 status=private
```

Restrictions:

- You cannot change the status of a volume from private to scratch if the volume belongs to a storage pool or is defined in the volume history file.
- Private volumes must be administrator-defined volumes with either no data or invalid data. They cannot be partially written volumes that contain active data. Volume statistics are lost when volume statuses are modified.

Removing volumes from an automated library

You can remove volumes from an automated library if you exported data to a volume and want to import the data to another system. You might also want to remove volumes to create space for new volumes.

About this task

By default, the server mounts the volume that you check out and verifies the internal label. When the label is verified, the server removes the volume from the library volume inventory, and then moves it to the entry/exit port or convenience I/O station of the library. If the library does not have an entry/exit port, the server requests that the mount operator remove the volume from a slot or device within the library.

Procedure

- To remove a volume from an automated library, issue the **CHECKOUT LIBVOLUME** command.
- For automated libraries with multiple entry/exit ports, issue the **CHECKOUT LIBVOLUME** command and specify the **REMOVE=BULK** parameter. The server ejects the volume to the next available entry/exit port.

What to do next

If you check out a volume that is defined in a storage pool and the server must access the volume later, the server requests that the volume be checked in. To return volumes to a library, issue the **CHECKIN LIBVOLUME** command.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[CHECKOUT LIBVOLUME \(Check a storage volume out of a library\)](#)

Maintaining a supply of scratch volumes in an automated library

When you define a storage pool that is associated with an automated library, you can specify a maximum number of scratch volumes equal to the physical capacity of the library. If the server is using a greater number of scratch volumes for the storage pool, you must ensure that enough volumes are available.

Procedure

If the number of scratch volumes that the server is using for the storage pool exceeds the number that is specified in the storage pool definition, complete the following steps:

1. Add scratch volumes to the library by issuing the **CHECKIN LIBVOLUME** command.

Tip: You might have to use an overflow location to move volumes out of the library to make room for these scratch volumes. For more information, see [“Managing a full library with an overflow location”](#) on page 186.

2. Increase the maximum number of scratch volumes that can be added to a storage pool by issuing the **UPDATE STGPOOL** command and specifying the **MAXSCRATCH** parameter.

What to do next

You might need more volumes for future recovery operations, so consider labeling and setting aside extra scratch volumes.

Related tasks

[Maintaining a supply of scratch volumes](#)

You must set the maximum number of scratch volumes for a storage pool high enough for the expected usage.

Managing a full library with an overflow location

As the demand for storage grows, the number of volumes that you need for a storage pool might exceed the physical capacity of an automated library. To make space available for new volumes and to monitor existing volumes, you can define an overflow location for a storage pool.

About this task

The server tracks the volumes that are moved to the overflow area and makes storage slots available for new volumes.

Procedure

1. Create a volume overflow location. Define or update the storage pool that is associated with the automated library by issuing the **DEFINE STGPOOL** or **UPDATE STGPOOL** command and specifying the **OVFLOCATION** parameter.
For example, to create an overflow location that is named ROOM2948 for a storage pool that is named ARCHIVEPOOL, issue the following command:

```
update stgpool archivepool ovflocation=Room2948
```

2. When you need to create space in the library for scratch volumes, move full volumes to the overflow location by issuing the **MOVE MEDIA** command.
For example, to move all full volumes in the specified storage pool out of the library, issue the following command:

```
move media * stgpool=archivepool
```

3. Check in scratch volumes as needed.

Restriction: If a volume has an entry in the volume history file, you cannot check it in as a scratch volume. For more information, see [“Checking volumes into an automated library”](#) on page 176.

4. Identify the empty scratch tapes in the overflow location by issuing the **QUERY MEDIA** command.
For example, issue the following command:

```
query media * stg=* whereovflocation=Room2948 wherestatus=empty
```

5. If the server requests additional volumes, locate and check in volumes from the overflow location.

To find volumes in an overflow location, issue the **QUERY MEDIA** command. You can also use the **QUERY MEDIA** command to generate commands by checking in volumes.

For example, to list the volumes in the overflow location, and at the same time generate the commands to check those volumes into the library, issue a command that is similar to the following example:

```
query media format=cmd stgpool=archivepool whereovflocation=Room2948
cmd="checkin libvol autolib &vol status=private"
cmdfilename="\storage\move\media\checkin.vols"
```

Tips:

- Mount requests from the server include the location of the volumes.
- To specify the number of days that must elapse before the volumes are eligible for processing, issue the **UPDATE STGPOOL** command and specify the **REUSEDELAY** parameter.
- The file that contains the generated commands can be run by using the IBM Storage Protect **MACRO** command.

Related information

[MOVE MEDIA \(Move sequential-access storage pool media\)](#)

[QUERY MEDIA \(Query sequential-access storage pool media\)](#)

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

Auditing the volume inventory in a library

You can audit an automated library to ensure that the library volume inventory is consistent with the volumes that are physically in the library. You might want to audit a library if the library volume inventory is distorted due to manual movement of volumes in the library or to database problems.

Procedure

1. Ensure that no volumes are mounted in the library drives. If any volumes are mounted in the IDLE state, issue the **DISMOUNT VOLUME** command to dismount them.
2. Audit the volume inventory by issuing the **AUDIT LIBRARY** command. Take one of the following actions:
 - If the library has a bar code reader, you can save time by using the bar code reader to identify volumes. For example, to audit the TAPELIB library by using its bar code reader, issue the following command:

```
audit library tapelib checklabel=barcode
```

- If the library does not have a bar code reader, issue the **AUDIT LIBRARY** command without specifying **CHECKLABEL=BARCODE**. The server mounts each volume to verify the label. After the label is verified, the server completes auditing any remaining volumes.

Results

The server deletes missing volumes from the inventory and updates the locations of volumes that moved since the last audit.

Restriction: The server cannot add new volumes to the inventory during an audit operation.

Related tasks

[Labeling tape volumes](#)

You must label tape volumes before the server can use them.

Related information

[AUDIT LIBRARY \(Audit volume inventories in an automated library\)](#)

[DISMOUNT VOLUME \(Dismount a volume by volume name\)](#)

Partially written volumes

Partially written volumes are always private volumes, even if their status was scratch before the server mounted them. The server tracks the original status of scratch volumes and returns them to scratch status when they are empty.

Except for volumes in automated libraries, the server is unaware of a scratch volume until after the volume is mounted. Then, the volume status changes to private, and the volume is automatically defined as part of the storage pool for which the mount request was made.

Related tasks

[Changing the status of a volume in an automated library](#)

You can change the status of a volume from private to scratch or from scratch to private.

Operations with shared libraries

Shared libraries are logical libraries that are represented physically by SCSI libraries. The physical library is controlled by the IBM Storage Protect server that is configured as a library manager. IBM Storage Protect servers that use the SHARED library type are library clients to the IBM Storage Protect library manager server.

The library client contacts the library manager when the library manager starts and the storage device initializes, or after a library manager is defined to a library client. The library client confirms that the contacted server is the library manager for the named library device. The library client also compares drive definitions with the library manager for consistency. The library client contacts the library manager for each of the following operations:

Volume mount

A library client sends a request to the library manager for access to a particular volume in the shared library device. For a scratch volume, the library client does not specify a volume name. If the library manager cannot access the requested volume, or if scratch volumes are unavailable, the library manager denies the mount request. If the mount is successful, the library manager returns the name of the drive where the volume is mounted.

Volume release

When a library client no longer needs to access a volume, it notifies the library manager that the volume can be returned to a scratch volume. The library manager database is updated with the new location for the volume, which is now in the inventory of the library server. The volume is deleted from the volume inventory of the library client.

Table 35 on page 188 shows the interaction between library clients and the library manager in processing IBM Storage Protect operations.

Table 35. How SAN-enabled servers process IBM Storage Protect operations		
Operation (Command)	Library manager	Library client
Query library volumes (QUERY LIBVOLUME)	Displays the volumes that are checked into the library. For private volumes, the owner server is also displayed.	Not applicable.
Check in and check out library volumes (CHECKIN LIBVOLUME , CHECKOUT LIBVOLUME)	Sends the commands to the library device.	Not applicable. When a check-in operation is required because of a client restore operation, a request is sent to the library manager server.

Table 35. How SAN-enabled servers process IBM Storage Protect operations (continued)

Operation (Command)	Library manager	Library client
Move media and move DRM media (MOVE MEDIA , MOVE DRMEDIA)	Valid only for volumes that are used by the library manager server.	Requests that the library manager server completes the operation. Generates a check-out process on the library manager server.
Audit library inventory (AUDIT LIBRARY)	Synchronizes the inventory with the library device.	Synchronizes the inventory with the library manager server.
Label a library volume (LABEL LIBVOLUME)	Labels and checks in volumes.	Not applicable.
Dismount a volume (DISMOUNT VOLUME)	Sends the request to the library device.	Requests that the library manager server completes the operation.
Query a volume (QUERY VOLUME)	Checks whether the volume is owned by the requesting library client and checks whether the volume is in the library device.	Requests that the library manager server completes the operation.

Managing server requests for volumes

IBM Storage Protect displays requests and status messages to all administrative command-line clients that are started in console mode. These request messages often have a time limit. Successful server operations must be completed within the time limit that is specified; otherwise, the operation times out.

About this task

For automated libraries, use the **CHECKIN LIBVOLUME** and **LABEL LIBVOLUME** commands to insert cartridges into slots. If you specify a value for the **WAITTIME** parameter, a reply message is displayed. If the value of the parameter is 0, no reply is required. When you issue the **CHECKOUT LIBVOLUME** command, you must insert cartridges into slots and, in all cases, a reply message is displayed.

Procedure

- The following table provides information about how to handle different server media tasks.

Task	Details
Use the administrative client for mount messages	The server sends mount request status messages to the server console and to all administrative command-line clients in mount mode or console mode. To start an administrative command-line client in mount mode, issue the dsmadmc -mountmode command on the administrative command-line client.
Receive messages about automated libraries	You can view mount messages and error messages about automated libraries on administrative command-line clients in mount mode or console mode. Mount messages are sent to the library and not to an operator. Messages about problems with the library are sent to the mount message queue.

Task	Details
Get information about pending operator requests	To get information about pending operator requests, issue the QUERY REQUEST command or view the mount message queue on an administrative command-line client that is started in mount mode. When you issue the QUERY REQUEST command, the server displays requested actions and the amount of time that is remaining before the requests time out.
Reply to operator requests	When the server requires an explicit reply to a completed mount request, use the REPLY command. The <i>request_number</i> parameter specifies the request identification number that tells the server which pending operator request is completed. This three-digit number is always displayed as part of the request message.
Cancel an operator request	To cancel a mount request for a library, issue the CANCEL REQUEST command. For most requests that are associated with automated SCSI libraries, an operator must complete a hardware or system action to cancel the requested mount. For such requests, the CANCEL REQUEST command is not accepted by the server. The CANCEL REQUEST command must include the request identification number. This number is included in the request message. If you want to mark the requested volume as UNAVAILABLE , issue the CANCEL REQUEST command and specify the PERMANENT parameter. If you specify the PERMANENT parameter, the server does not try to mount the requested volume again. This is useful if, for example, the volume is at a remote site or is otherwise unavailable.
Respond to a volume check-in request	If the server cannot find a particular volume to mount in an automated library, the server requests that the operator check in the volume. If the requested volume is available, place the volume in the library and check it in. For more information, see “Checking volumes into an automated library” on page 176 . If the requested volume is unavailable, update the access mode of the volume by issuing the UPDATE VOLUME command and specifying the ACCESS=UNAVAILABLE parameter. Then, cancel the check-in request by using the CANCEL REQUEST command. Do not cancel the client process that caused the request. Use the QUERY REQUEST command to obtain the ID of the request that you want to cancel. If you do not respond to the check-in request from the server within the mount-wait period that is specified for the device class for the storage pool, the server marks the volume as unavailable.
Determine which volumes are mounted	For a report about all volumes that are currently mounted for use by the server, issue the QUERY MOUNT command. The report shows which volumes are mounted, which drives accessed them, and whether the volumes are in use.
Dismount idle volumes	When a volume is idle, the server keeps it mounted for a time that is specified by the mount retention parameter for the device class. Using a mount retention value can reduce the access time when volumes are used repeatedly. To dismount an idle volume from the drive where it is mounted, issue the DISMOUNT VOLUME command. For information about setting mount retention times, see “Controlling the amount of time that a volume remains mounted” on page 126 .

Related information

[QUERY REQUEST](#) (Query one or more pending mount requests)

Managing tape drives

You can query, update, and delete tape drives. You can also clean tape drives and configure tape drive encryption and data validation.

Updating drives

You can change the attributes of a drive definition to take a drive offline or reconfigure it.

About this task

You can change the following attributes of a drive:

- The element address, if the drive is in a SCSI
- The cleaning frequency
- The drive status: online or offline

Restriction: If a drive is in use, you cannot change the element number or the device name. For instructions about taking drives offline, see [“Taking tape drives offline”](#) on page 192.

If a volume is mounted in the drive but the volume is idle, it can be explicitly dismounted. For instructions about dismounting idle volumes, see [“Managing server requests for volumes”](#) on page 189.

Procedure

- Change the element address of a drive by issuing the **UPDATE DRIVE** command.
For example, in a library that is named AUTO, change the element address of DRIVE3 to 119 by issuing the following command:

```
update drive auto drive3 element=119
```

- Change the device name of a drive by issuing the **UPDATE PATH** command.
For example, to change the device name of a drive that is named DRIVE3, issue the following command:

```
AIX update path server1 drive3 srctype=server desttype=drive library=scsilib  
device=/dev/rmt0
```

```
Linux update path server1 drive3 srctype=server desttype=drive library=scsilib  
device=/dev/IBMtape0
```

```
Windows update path server1 drive3 srctype=server desttype=drive library=scsilib  
device=mt3.0.0.0
```

Related information

[UPDATE DRIVE](#) (Update a drive)

[UPDATE PATH](#) (Change a path)

Taking tape drives offline

You can take a tape drive offline while it is in use. For example, you might take a drive offline to complete maintenance.

About this task

If you change the status of a drive to offline while the drive is in use, the server finishes processing the tape that is in the drive, and then stops using the drive. However, if the tape that was in use was part of a series of tapes for a single transaction, the drive is unavailable to complete the series. If no other drives are available, the transaction might fail.

Procedure

- To change the status of a drive, issue the **UPDATE DRIVE** command and specify the **ONLINE** parameter. For example, to update the DRIVE3 drive in the MANLIB library and take the drive offline, issue the following command:

```
update drive manlib drive3 online=no
```

Restriction: Do not specify other optional parameters when you specify the **ONLINE** parameter. If you do, the drive is not updated, and the command fails when the drive is in use.

Results

If you update all drives in a library to an offline status, processes that require a library mount point fail.

The updated state of the drive is retained even when the server is halted and restarted. If a drive is marked offline when the server is restarted, a warning is issued stating that the drive must be manually brought online.

Related information

[UPDATE DRIVE \(Update a drive\)](#)

Data validation during read/write operations to tape

To validate data and identify data that is corrupted, you can use a feature that is called logical block protection. If you use logical block protection, IBM Storage Protect inserts a cyclic redundancy check (CRC) value at the end of each logical block of data while it is written to tape.

With logical block protection, you can identify errors that occur when data is written to tape and during data transfer from the tape drive to IBM Storage Protect through the storage area network. Drives that support logical block protection validate data during read and write operations. The IBM Storage Protect server validates data during read operations.

If validation by the drive fails during write operations, the failure can indicate that data was corrupted during transfer to tape. In this case, the IBM Storage Protect server fails the write operation. You must restart the operation to continue. If validation by the drive fails during read operations, the failure can indicate that the tape media is corrupted. If validation by the IBM Storage Protect server fails during read operations, the failure can indicate that data was corrupted during transfer from the tape drive, and the server tries the operation again. If validation fails consistently, the IBM Storage Protect server issues an error message that indicates hardware or connection problems.

If logical block protection is disabled on a tape drive, or the drive does not support logical block protection, the IBM Storage Protect server can read protected data. However, the data is not validated.

Logical block protection is superior to the CRC validation that you can specify when you define or update a storage pool. When you specify CRC validation for a storage pool, data is validated only during volume auditing operations. Errors are identified after the data is written to tape.

Restrictions:

- You cannot use logical block protection for sequential data such as backup sets and database backups.
- CRC checking impacts performance because more processor usage is required on both the client and server to calculate and compare CRC values.
- For a scratch volume, if you specify logical block protection for read/write operations (**LBPROTECT=READWRITE**), do not change the parameter value at any time after data is written to the volume. Changing the parameter value during the life of the volume on the IBM Storage Protect server is not supported.

Drives that support logical block protection

Logical block protection is available only for 3592, LTO, and ECARTRIDGE device types. Capable 3592 drives include IBM TS1130, TS1140, and later generations. Capable LTO drives include IBM LTO-5 and supported LTO-6 drives. Capable Oracle StorageTek drives include drives with the T10000C and T10000D format.

The following table shows the media and the formats that you can use with drives that support logical block protection.

Drive	Tape media	Drive formats
IBM TS1130	3592 Generation 2	3592-3 and 3592-3C
IBM TS1140	3592 Generation 2 3592 Generation 3	Generation 2: 3592-3 and 3592-3C Generation 3: 3592-4 and 3592-4C
IBM TS1150	3592 Generation 3 3592 Generation 4	Generation 4: 3592-5 and 3592-5C
IBM LTO-5	LTO-5	Ultrium 5 and Ultrium 5C
IBM LTO-6	LTO-6 LTO-5	Ultrium 6 and Ultrium 6C Ultrium 5 and Ultrium 5C
IBM LTO-7	LTO-7 LTO-6	Ultrium 7 and Ultrium 7C Ultrium 6 and Ultrium 6C
Oracle T10000C	Oracle StorageTek T10000 T2	T10000C and T10000C-C
Oracle T10000D	Oracle StorageTek T10000 T2	T10000D and T10000D-C

Tips:

- To enable logical block protection for a tape volume and then reuse the volume to back up data, you must enable logical block protection for the device class and the drive.
- If you have a 3592, LTO, or Oracle StorageTek drive that is not capable of logical block protection, you can upgrade the drive with firmware that provides logical block protection.

Logical block protection is available for drives that are in SCSI libraries. For the most current information about support for logical block protection, see [technote 1568108](#).

To use logical block protection for write operations, all drives in the library must support logical block protection. If a drive is not capable of logical block protection, volumes that have read/write access are not mounted. However, the server can use the drive to mount volumes that have read-only access. The protected data is read and validated by the IBM Storage Protect server if logical block protection is enabled for read/write operations.

Enabling and disabling logical block protection

You can specify logical block protection for read and write operations, or only for write operations. You can also disable logical block protection. By default, logical block protection is disabled because of performance effects that result from cyclic redundancy check (CRC) validation on the server and the tape drive.

About this task

Read/write operations to empty or filling volumes depend on whether the volumes have logical block protection. Protected and unprotected data blocks cannot be mixed on the same volume. If you change the setting for logical block protection, the change applies only to empty volumes. Filling and full volumes maintain their status of logical block protection until they are empty and ready to be refilled. For example, if you disable logical block protection and the server selects a volume that is associated with a device class that has logical block protection, the server continues writing protected data to the volume.

Restriction: Logical block protection is available only for certain device types. For more information, see [“Drives that support logical block protection”](#) on page 193.

Procedure

1. To enable logical block protection for the 3592, LTO, and ECARTRIDGE device types, issue the **DEFINE DEVCLASS** or the **UPDATE DEVCLASS** command and specify the **LBPROTECT** parameter. For example, to specify logical block protection during read and write operations for a 3592 device class that is named 3592_lbprotect, issue the following command:

```
define devclass 3592_lbprotect library=3594 lbprotect=readwrite
```

Tips:

- If you update the value of the **LBPROTECT** parameter from NO to READWRITE or WRITEONLY and the server selects a filling volume without logical block protection for write operations, the server issues a message each time the volume is mounted. The message indicates that data is written to the volume without logical block protection. To prevent this message from displaying or to have IBM Storage Protect write data only with logical block protection, update the access of filling volumes without logical block protection to read-only.
 - To improve performance, do not specify the **CRCDATA** parameter on the **DEFINE STGPOOL** or **UPDATE STGPOOL** command.
 - When data is validated during read operations by both the drive and by the IBM Storage Protect server, it can slow server performance during restore and retrieve operations. To reduce the time that is required for restore and retrieve operations, change the setting of the **LBPROTECT** parameter from READWRITE to WRITEONLY. After data is restored or retrieved, you can reset the **LBPROTECT** parameter to READWRITE.
2. To disable logical block protection, issue the **DEFINE DEVCLASS** or the **UPDATE DEVCLASS** command and specify the **LBPROTECT=NO** parameter.

Restriction: If logical block protection is disabled, the server does not write to an empty tape with logical block protection. However, if a filling volume with logical block protection is selected, the server continues to write to the volume with logical block protection. To prevent the server from writing to tapes with logical block protection, change the access of filling volumes with logical block protection to read-only. When data is read, the CRC results are not checked by the drive or server.

If a disaster occurs and the disaster recovery site does not have drives that support logical block protection, you must specify the **LBPROTECT=NO** parameter. If the tape drives are used for write operations, you must change the volume access for volumes with protected data to read-only to prevent the server from using the volumes.

If the server must enable logical block protection, the server issues an error message that indicates that the drive does not support logical block protection.

What to do next

To determine whether a volume has logical block protection, issue the **QUERY VOLUME** command and review the value in the Logical Block Protection field.

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

[QUERY VOLUME \(Query storage pool volumes\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

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

Read/write operations to volumes with logical block protection

Read/write operations to empty or filling volumes depend on whether the volumes have logical block protection. Protected and unprotected data blocks cannot be mixed on the same volume.

If you use the **UPDATE DEVCLASS** command to change the setting for logical block protection, the change applies only to empty volumes. Filling and full volumes maintain their status of logical block protection until they are empty and ready to be refilled.

For example, suppose that you change the value of the **LBPROTECT** parameter from READWRITE to NO. If the server selects a volume that is associated with the device class and that has logical block protection, the server continues writing protected data to the volume.

Tips:

- If a drive does not support logical block protection, volumes with logical block protection for write operations cannot be mounted. To prevent the server from mounting the protected volumes for write operations, change the volume access to read-only. Also, disable logical block protection to prevent the server from enabling the feature on the tape drive.
- If a drive does not support logical block protection, and logical block protection is disabled, the server reads data from protected volumes. However, the data is not validated by the server and the tape drive.

Related information

[QUERY VOLUME \(Query storage pool volumes\)](#)

[UPDATE DEVCLASS \(Update a device class\)](#)

Storage pool management in a tape library

To mix protected and unprotected data in a library, you must create different device classes and different storage pools to separate the data. If a device class is associated with protected data, you can specify logical block protection for read and write operations or for write operations only.

To define device classes and storage pools for a TS3500 library that has LTO-5 drives, for protected and unprotected data, you can issue a series of commands as shown in the following example:

```
define library 3584 libtype=scsi
define devclass lbprotect library=3584 devicetype=lto lbprotect=readwrite
define devclass normal library=3584 devicetype=lto lbprotect=no
define stgpool lbprotect_pool lbprotect maxscratch=10
define stgpool normal_pool normal maxscratch=10
```

Related information

[DEFINE DEVCLASS \(Define a device class\)](#)

[DEFINE LIBRARY \(Define a library\)](#)

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

Cleaning tape drives

You can use the server to manage tape-drive cleaning. The server can control how tape drives in SCSI libraries are cleaned.

About this task

You must have system privilege or unrestricted storage privilege to clean tape drives. For automated libraries, you can automate cleaning by specifying the frequency of cleaning operations and checking a cleaner cartridge into the library volume inventory. IBM Storage Protect mounts the cleaner cartridge as specified. There are special considerations if you plan to use server-controlled drive cleaning with a SCSI library that provides automatic drive cleaning support in its device hardware.

Tip: If an automated tape library supports library-drive cleaning, ensure that the feature is enabled.

You can prevent premature wear on the read/write heads of drives by using the library cleaning functions that are available from your device manufacturer.

Drives and libraries from manufacturers differ in how they manage cleaner cartridges, and how they report the presence of a cleaner cartridge in a drive. The device driver might not be able to open a drive that contains a cleaner cartridge. Sense codes and error codes that are issued by devices for drive cleaning vary. Library-drive cleaning is usually not known to applications. Therefore, IBM Storage Protect might not always detect the cleaner cartridges in drives and might not be able to determine when cleaning begins.

Some devices require a small amount of idle time between mount requests to start drive cleaning. However, IBM Storage Protect tries to minimize the idle time for a drive. The result might be to prevent the library drive cleaning from functioning effectively. If this happens, use IBM Storage Protect to control drive cleaning. You can set the frequency to match the cleaning recommendations from the manufacturer.

Methods for cleaning tape drives

Over time, the read heads on tapes can get dirty, which can cause read and write operations to fail. To prevent these issues, enable tape cleaning. You can enable tape cleaning from the drive or from IBM Storage Protect.

You can choose to use the library-drive cleaning method or the IBM Storage Protect drive-cleaning method, but not both. Some SCSI libraries provide automatic drive cleaning. Select the library-drive cleaning method if it is available. If it is unavailable or causes issues, use IBM Storage Protect to control library drive cleaning.

Library drive-cleaning method

The library drive-cleaning method provides several advantages for automated tape libraries that use this function:

- Reduces the burden on the IBM Storage Protect administrator to physically manage cartridge cleaning.
- Improves cleaning cartridge usage rates. Most tape libraries track the number of times that drives can be cleaned based on hardware indicators. IBM Storage Protect uses a raw count.
- Reduces unnecessary cleaning. Modern tape drives do not have to be cleaned at fixed intervals and can detect and request when cleaning is required.

Manufacturers who provide a library drive-cleaning method recommend its use to prevent premature wear on the read/write heads of the drives. Drives and libraries from different manufacturers differ in how they manage cleaner cartridges and how they report the presence of a cleaner cartridge in a drive. The device driver might not be able to open a drive that contains a cleaner cartridge. Sense codes and error codes that are issued by devices for drive cleaning vary. Library drive cleaning is usually transparent to all applications. However, IBM Storage Protect might not always detect cleaner cartridges in drives and might not be able to determine when cleaning begins.

IBM Storage Protect drive cleaning method

Some devices require a small amount of idle time between mount requests to start drive cleaning. However, IBM Storage Protect tries to minimize the idle time for a drive. The result might be to prevent the library drive cleaning from functioning effectively. If this happens, try using IBM Storage Protect to control drive cleaning. Set the frequency to match the cleaning recommendations from the manufacturer.

If IBM Storage Protect controls the drive-cleaning process, disable the library drive-cleaning function to prevent problems. If the library drive-cleaning function is enabled, some devices automatically move any cleaner cartridge that is found in the library to slots in the library that are dedicated to cleaner cartridges. You cannot check a cleaner cartridge into the IBM Storage Protect library inventory until you disable the library drive-cleaning function.

To enable cleaning from the drive, follow the instructions that are provided by the drive manufacturer. To enable cleaning by using IBM Storage Protect, see [“Configuring the server for drive cleaning in an automated library”](#) on page 197.

Configuring the server for drive cleaning in an automated library

When you configure server-controlled drive cleaning in an automated library, you can specify how often you want the drives to be cleaned.

Before you begin

Determine how often the drive must be cleaned. This step is required so that you can specify an appropriate value for the **CLEANFREQUENCY** parameter on the **DEFINE DRIVE** or **UPDATE DRIVE** command. For example, to clean a drive after 100 GB of data is processed on the drive, you would specify **CLEANFREQUENCY=100**.

For guidelines about cleaning frequency, see the drive manufacturer's documentation. If the documentation provides guidelines for cleaning frequency in terms of hours of use, convert the value to a gigabyte value by completing the following steps:

1. Use the bytes-per-second value for the drive to determine a gigabytes-per-hour value.
2. Multiply the gigabytes-per-hour value by the recommended hours of use between cleanings.
3. Use the result as the cleaning frequency value.

You can either specify a value for the **CLEANFREQUENCY** parameter or specify **ASNEEDED** to clean the drive as needed.

Restrictions:

1. For IBM 3592 drives, you must specify a numerical value for the **CLEANFREQUENCY** parameter. By using the cleaning frequency that is listed in the product documentation, you will not overclean the drives.
2. The **CLEANFREQUENCY=ASNEEDED** parameter value does not work for all tape drives. To determine whether a drive supports this function, see the information for your operating system:

AIX	Windows	Supported devices for AIX and Windows
Linux		Supported devices for Linux

In the technote, click the drive name to view detailed information. If the ASNEEDED value is not supported, specify the number of gigabytes.

Procedure

Define or update the drives in the library, by using the **CLEANFREQUENCY** parameter in the **DEFINE DRIVE** or **UPDATE DRIVE** command.

For example, to clean a drive that is named DRIVE1 after 100 GB of data is processed, issue the following command:

```
update drive autolib1 drive1 cleanfrequency=100
```

Results

After the cleaner cartridge is checked in, the server mounts the cleaner cartridge in a drive when the drive needs cleaning. The server uses that cleaner cartridge for the number of specified cleanings. For more information, see [“Operations with cleaner cartridges” on page 152](#).

What to do next

Check the cleaner cartridge into the library volume inventory by following the instructions in [“Checking a cleaner cartridge into a library” on page 198](#).

Related information

[DEFINE DRIVE \(Define a drive to a library\)](#)

[UPDATE DRIVE \(Update a drive\)](#)

Checking a cleaner cartridge into a library

To enable automatic tape-drive cleaning, you must check a cleaner cartridge into the volume inventory of the automated library.

About this task

When you check a cleaner cartridge into a library, ensure that it is correctly identified to the server as a cleaner cartridge. Ensure that a cleaner cartridge is not in a slot that is detected by the search process. Errors and delays of 15 minutes or more might indicate that a cleaner cartridge is improperly placed.

The preferred method is to check in cleaner cartridges individually. If you have to check in both data cartridges and cleaner cartridges, place the data cartridges in the library and check them in first. Then, check the cleaner cartridge in to the library.

Procedure

To check a cleaner cartridge into a library, issue the **CHECKIN LIBVOLUME** command.

For example, to check in a cleaner cartridge that is named AUTOLIB1, issue the following command:

```
checkin libvolume autolib1 cleanv status=cleaner cleanings=10  
checklabel=no
```

The server requests that the cartridge is placed in the entry/exit port, or into a specific slot.

Related information

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

Operations with cleaner cartridges

To ensure that tape drives are cleaned when necessary, and to avoid issues with tape storage, follow the guidelines.

Monitoring the cleaning process

If a cleaner cartridge is checked in to a library, and a drive must be cleaned, the server dismounts the data volume and runs the cleaning operation. If the cleaning operation fails or is canceled, or if no cleaner cartridge is available, you might not be aware that the drive needs cleaning. Monitor cleaning messages for these problems to ensure that drives are cleaned as needed. If necessary, issue the **CLEAN DRIVE** command to have the server try the cleaning again, or manually load a cleaner cartridge into the drive.

Using multiple cleaner cartridges

The server uses a cleaner cartridge for the number of cleanings that you specify when you check in the cleaner cartridge. If you check in two or more cleaner cartridges, the server uses only one of the cartridges until the designated number of cleanings for that cartridge is reached. Then, the server

uses the next cleaner cartridge. If you check in two or more cleaner cartridges and issue two or more **CLEAN DRIVE** commands concurrently, the server uses multiple cartridges at the same time and decrements the remaining cleanings on each cartridge.

Related information

[AUDIT LIBRARY \(Audit volume inventories in an automated library\)](#)

[CHECKIN LIBVOLUME \(Check a storage volume into a library\)](#)

[CLEAN DRIVE \(Clean a drive\)](#)

[LABEL LIBVOLUME \(Label a library volume\)](#)

[QUERY LIBVOLUME \(Query a library volume\)](#)

Resolving errors that are related to drive cleaning

While moving cartridges within a library, you might place a data cartridge where a cleaner cartridge should be. Review the process that the server completes and the messages that are issued so that you can resolve the issue.

When a drive needs cleaning, the server loads what its database shows as a cleaner cartridge into the drive. The drive then moves to a READY state, and IBM Storage Protect detects that the cartridge is a data cartridge. The server completes the following steps:

1. The server attempts to read the internal tape label of the data cartridge.
2. The server ejects the cartridge from the drive and moves it back to the home slot of the cleaner cartridge within the library. If the eject operation fails, the server marks the drive offline and issues a message that the cartridge is still in the drive.
3. The server checks out the cleaner cartridge to avoid selecting it for another drive cleaning request. The cleaner cartridge remains in the library but no longer appears in the IBM Storage Protect library inventory.
4. By using the internal tape label, the server checks the volume name against the current library inventory, storage pool volumes, and the volume history file.
 - If the volume name is not found in the library inventory, a data cartridge might be checked in as a cleaner cartridge by mistake. When the volume is checked out, you do not have to take further action.
 - If the volume name is found in the library inventory, the server issues messages that manual intervention and a library audit are required. To resolve the issue, follow the instructions in [“Auditing the volume inventory in a library”](#) on page 187.

Tape drive replacement

If you replace a drive in a tape library that is defined to IBM Storage Protect, you must delete the drive and path definitions for the old drive and define the new drive and path.

Replacing drive and path definitions is required even if you are exchanging one drive for another of the same type, with the same logical address, physical address, SCSI ID, and port number. Device alias names can change when you change your drive connections.

If the new drive is an upgrade that supports a new media format, you might be required to define a new logical library, device class, and storage pool. Procedures for setting up a policy for a new drive in a multiple-drive library vary, depending on the types of drives and media in the library.

Deleting tape drives

You can delete tape drives from a library. For example, you can delete a drive that you no longer use, or a drive that you want to replace.

Procedure

1. Stop the IBM Storage Protect server and shut down the operating system.

2. Remove the old drive and follow the manufacturer's instructions to install the new drive.
3. Restart the operating system and the IBM Storage Protect server.
4. Delete the path from the server to the drive.

For example, to delete a path from SERVER1 to LIB1, issue the following command:

```
delete path server1 lib1 srctype=server desttype=drive
```

5. Delete the drive definition.

For example, issue the following command to delete a drive that is named DLT1 from a library device that is named LIB1:

```
delete drive lib1 dlt1
```

Related information

[DELETE DRIVE \(Delete a drive from a library\)](#)

[DELETE PATH \(Delete a path\)](#)

Replacing drives with others of the same type

To add a drive that supports the same media formats as the drive it replaces, you must define a new drive and path.

About this task

If a library includes only one model of drive and you want to replace a drive, you must replace the drive with the same model drive. If a library includes mixed models of drives and you want to replace a drive, you can replace the drive with any model drive that exists in the library.

Procedure

1. Delete the path and drive definitions for the old drive. For example, to delete a drive that is named DRIVE1 from a library that is named LIB1, enter the following command:

```
delete path server2 drive1 srctype=server desttype=drive library=lib1
delete drive lib1 drive1
```

2. Power off the library, remove the original drive, replace it with the new drive, and power on the library.
3. Refresh the host system to ensure that the system detects the new drive.
4. Define the new drive and path. For example, to define a new drive, DRIVE2, and a path to it from SERVER2, if you are using the IBM Storage Protect device driver, enter the following commands:

```
AIX define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=/dev/mt0
```

```
Linux define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=/dev/tmscsi/mt0
```

```
Windows define drive lib1 drive2
define path server2 drive2 srctype=server desttype=drive library=lib1
device=mt3.0.0.1
```

Tip: You can use your existing library, device class, and storage pool definitions.

Related information

[DELETE DRIVE \(Delete a drive from a library\)](#)

[DELETE PATH \(Delete a path\)](#)

Migrating data to upgraded drives

If you upgrade all of the tape drives in a library, you can preserve your existing policy definitions to migrate and expire existing data, and you can use the new drives to store data.

Before you begin

The following scenario assumes that you already have a primary storage pool for a DISK device class that is named POOL1.

Procedure

1. To migrate data to a storage pool that is created for the new drives, specify the **NEXTSTGPOOL** parameter. For example, to migrate data from an existing storage pool, POOL1, to the new storage pool, POOL2, issue the following command:

```
update stgpool pool1 nextstgpool=pool2
```

2. Update the management-class definitions to store data in the DISK storage pool by using the **UPDATE MGMTCLASS** command.

Related information

[DEFINE STGPOOL \(Define a volume in a storage pool\)](#)

[UPDATE MGMTCLASS \(Update a management class\)](#)

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

Securing the IBM Storage Protect 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.

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">a. Register the administrator and specify Pa\$#\$tw0 as the password by issuing the following command: <pre>register admin admin1 Pa\$#\$tw0</pre>b. Grant system authority to the administrator by issuing the following command: <pre>grant authority admin1 classes=system</pre>

Task	Procedure
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:</p> <pre>remove admin admin1</pre>
Temporarily prevent access to the server.	Lock or unlock an administrator by using the LOCK ADMIN or UNLOCK ADMIN command.

Related concepts

[Planning for administrator roles](#)

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

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 36. 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.

Table 36. Authentication tasks for IBM Storage Protect servers (continued)

Task	Procedure
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>
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>

Table 36. Authentication tasks for IBM Storage Protect servers (continued)

Task	Procedure
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.

Securing the server 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

1. After you register an administrator by using the **REGISTER ADMIN** command, use the **GRANT AUTHORITY** command to set the administrator's authority level.
For details about setting and changing authority, see [“Managing administrators” on page 201](#).
2. To control the authority of an administrator to complete some tasks, use the following two server options:
 - a) You can select the authority level that an administrator must have to issue **QUERY** and **SELECT** commands with the **QUERYAUTH** server option. By default, no authority level is required. You can change the requirement to one of the authority levels, including system.
 - b) You can specify that system authority is required for commands that cause the server to write to an external file with the **REQSYSAUTHOUTFILE** server option. By default, system authority is required for such commands.
3. You can restrict data backup on a client node to only root user IDs or authorized users.
For example, to limit backups to the root user ID, issue the **REGISTER NODE** or **UPDATE NODE** command and specify the **BACKUPINITIATION=root** parameter:

```
update node backupinitiation=root
```

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 period

If 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, `dsmserve.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:

```
dsmserve 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. For information about scheduling activities such as backing up the database, backing up the device configuration file, and backing up the volume history, see [“Defining schedules for server maintenance activities” on page 58](#).

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**.

Related information

[HALT \(Shut down the server\)](#)

Preparing for and recovering from a disaster by using DRM

IBM Storage Protect provides a disaster recovery manager (DRM) function to recover your server and client data during a disaster.

DRM tracks the movement of offsite media and registers that information in the IBM Storage Protect database. DRM consolidates plans, scripts, and other information in a plan file that is required to recover the IBM Storage Protect server when a disaster or unplanned outage occurs. DRM stores copies of backup files that are required to restore a server in the recovery plan file. The following backup files and directories are required to restore a server:

- Server options file (`dsmserve.opt`)
- Device configuration file (for example, `devconf.dat`)
- Volume history file (for example, `volhist.dat`)
- Master encryption key files (`dsmkeydb.kdb` or `dsmkeydb.sth`)
- Server certificate and private key files (`cert.kbd` or `cert.sth`)

If you are concerned about possible malware attacks, including ransomware, consider using DRM because it can help you recover your servers after an attack.

Restriction: DRM is available only in the IBM Storage Protect Extended Edition product.

Disaster recovery plan file

The disaster recovery plan file contains the information that is required to recover an IBM Storage Protect server to the point in time of the last database backup operation that was completed before the plan was created.

The plan is organized into stanzas, which you can separate into multiple files. Each stanza has a begin statement and an end statement.

Table 37. Stanzas in the disaster recovery plan file	
Stanza	Information in the stanza
SERVER.REQUIREMENTS	Identifies the database and recovery log storage requirements for the server.
RECOVERY.INSTRUCTIONS.GENERAL	Identifies site-specific instructions that the administrator enters in the file that is identified by the prefix <code>RECOVERY.INSTRUCTIONS.GENERAL</code> . The instructions include the recovery strategy, key contact names, an overview of key applications that are backed up by this server, and other relevant recovery instructions.
RECOVERY.INSTRUCTIONS.OFFSITE	Contains instructions that the administrator enters in the file that is identified by the prefix <code>RECOVERY.INSTRUCTIONS.OFFSITE</code> . The instructions describe the name and location of the offsite vault, and how to contact the vault administrator (for example, a name and phone number).
RECOVERY.INSTRUCTIONS.INSTALL	Contains instructions that the administrator enters in the file that is identified by the prefix <code>RECOVERY.INSTRUCTIONS.INSTALL</code> . The instructions describe how to rebuild the base server and provide the location of the system image backup copies.
RECOVERY.INSTRUCTIONS.DATABASE	Contains instructions that the administrator enters in the file that is identified by the prefix <code>RECOVERY.INSTRUCTIONS.DATABASE</code> . The instructions describe how to prepare for the database recovery. For example, you might enter instructions about how to initialize or load the backup volumes for an automated library. No sample of this stanza is provided.

Table 37. Stanzas in the disaster recovery plan file (continued)	
Stanza	Information in the stanza
RECOVERY.INSTRUCTIONS.STGPOOL	Contains instructions that the administrator enters in the file that is identified by the prefix RECOVERY.INSTRUCTIONS.STGPOOL. The instructions include the names of your software applications and the copy storage pool names that contain the backups of these applications. No sample of this stanza is provided.
RECOVERY.VOLUMES.REQUIRED	Provides a list of the database backup and copy storage pool volumes that are required to recover the server. A database backup volume is included if it is part of the most recent database backup series. A copy storage pool volume is included if it is not empty and not marked destroyed.
RECOVERY.DEVICES.REQUIRED	Provides details about the devices that are required to read the backup volumes.
RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE	Contains a script with the commands that are required to recover the server.
RECOVERY.SCRIPT.NORMAL.MODE	Contains a script with the commands that are required to restore the server primary storage pools.
DB.STORAGEPATHS	Identifies the directories for the IBM Storage Protect database.
LICENSE.REGISTRATION	Contains a macro to register your server licenses.
COPYSTGPOOL.VOLUMES.AVAILABLE	Contains a macro to mark copy storage pool volumes that were moved offsite and then moved back onsite. You can use the information as a guide and issue the administrative commands. Alternatively, copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
COPYSTGPOOL.VOLUMES.DESTROYED	Contains a macro to mark copy storage pool volumes as unavailable if the volumes were onsite at the time of the disaster. These volumes are considered offsite and have not been destroyed in a disaster. You can use the information as a guide and issue the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
PRIMARY.VOLUMES.DESTROYED	Contains a macro to mark primary storage pool volumes as destroyed if the volumes were onsite at the time of disaster. You can use the information as a guide and run the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
PRIMARY.VOLUMES.REPLACEMENT	Contains a macro to identify replacement primary storage pool volumes. You can use the information as a guide and run the administrative commands from a command line, or you can copy, modify, and run the macro to a file. This macro is started by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
STGPools.RESTORE	Contains a macro to restore the primary storage pools. You can use the stanza as a guide and run the administrative commands from a command line. You can also copy, modify, and run it to a file. This macro is started by the RECOVERY.SCRIPT.NORMAL.MODE script.
VOLUME.HISTORY.FILE	Contains a copy of the volume history information when the recovery plan was created. The DSMSERV RESTORE DB utility uses the volume history file to determine what volumes are needed to restore the database. The volume history file is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
DEVICE.CONFIGURATION.FILE	Contains a copy of the server device configuration information when the recovery plan was created. The DSMSERV RESTORE DB utility uses the device configuration file to read the database backup volumes. The device configuration file is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.
DSMSERV.OPT.FILE	Contains a copy of the server options file. This stanza is used by the RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script.

Table 37. Stanzas in the disaster recovery plan file (continued)	
Stanza	Information in the stanza
LICENSE.INFORMATION	Contains a copy of the latest license audit results and the server license terms.
MACHINE.GENERAL.INFORMATION	Provides information for the server machine, such as its location, which is needed to rebuild the server machine. This stanza is included in the plan file if the machine information is saved in the database by using the DEFINE MACHINE command and specifying the ADSMSEVER=YES .
MACHINE.RECOVERY.INSTRUCTIONS	Provides the recovery instructions about the server machine. This stanza is included in the plan file if the machine recovery instructions are saved in the database.
MACHINE.RECOVERY.CHARACTERISTICS	Provides the hardware and software characteristics for the server machine. This stanza is included in the plan file if the machine characteristics are saved in the database.
MACHINE.RECOVERY.MEDIA	Provides information about the media that are required for rebuilding the machine that contains the server. This stanza is included in the plan file if recovery media information is saved in the database and it is associated with the machine that contains the server.

Recovering the server and client data by using DRM

Use the disaster recovery manager (DRM) function to recover the IBM Storage Protect server and client data when a disaster occurs.

Before you begin

IBM Storage Protect is set up to use the Secure Sockets Layer (SSL) protocol for client/server authentication. When you start the server, a digital certificate file, `cert.kdb`, is created as part of the process. This file includes the server's public key, which allows the client to encrypt data. The digital certificate file cannot be stored in the server database because the Global Security Kit (GSKit) requires a separate file in a certain format.

1. Keep backup copies of the `cert.kdb`, `cert.sth`, and `cert256.arm` files.
2. If both the original certificate files and any copies are lost or corrupted, generate new certificate files.

The master encryption key is stored in a new GSKit-managed key database, `dsmkeydb.kdb`. If the server has an existing master encryption key, the master encryption key is migrated from the `dsmserv.pwd` file to the key database, `dsmkeydb.kdb`. Keep backup copies of the `dsmkeydb.kdb` and `dsmkeydb.sth` files. You can configure the **BACKUP DB** command to back up the master encryption key, or you can manually back up the `dsmkeydb.kdb` and `dsmkeydb.sth` files yourself. You cannot recover from a disaster without the master encryption key.

1. Keep backup copies of the `dsmkeydb.kdb` and `dsmkeydb.sth` files.

Procedure

1. Get the latest recovery plan.
2. Review the recovery steps that are described in the `RECOVERY.INSTRUCTIONS.GENERAL` stanza of the plan.
3. Separate the stanzas of the plan file into individual files for general preliminary instructions, IBM Storage Protect server recovery scripts, and client recovery instructions.
4. Retrieve all required recovery volumes (as listed in the plan) from the vault.
5. Review the device configuration file to ensure that the hardware configuration at the recovery site is the same as the original site. Any differences must be updated in the device configuration file. The following example configuration changes require updates to the configuration information:

- Different device names.
 - For automated libraries, the requirement of manually placing the database backup volumes in the automated library and updating the configuration information to identify the element within the library. This allows the server to locate the required database backup volumes.
6. Set up replacement hardware for the IBM Storage Protect server, including the operating system and the IBM Storage Protect base release installation.
 7. Run the IBM Storage Protect server recovery scripts from the recovery plan. The RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE and RECOVERY.SCRIPT.NORMAL.MODE stanzas contain executable command files that can be used to drive the recovery of the IBM Storage Protect server by calling other command files that were generated in the plan. The RECOVERY.SCRIPT.DISASTER.RECOVERY.MODE script recovers the server to the point where clients can begin restores directly from the copy storage pool volumes.
 8. Restore the primary storage pools by using the RECOVERY.SCRIPT.NORMAL.MODE script.
 9. Start client restore operations in order of highest priority, as defined in your high-level planning.

What to do next

The IBM Storage Protect server can now be used for normal server operations. Ensure that all required operations are scheduled. For instructions, see [“Defining schedules for server maintenance activities” on page 58](#) and [Scheduling backup and archive operations](#).

Related information

[PREPARE \(Create a recovery plan file\)](#)

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

Running a disaster recovery drill

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.

Before you begin

Complete the following tasks:

- Schedule activities regularly to manage, protect, and maintain the server. For more information about scheduling activities, see [“Defining schedules for server maintenance activities” on page 58](#). Ensure that you schedule the following tasks:
 - Backing up the database.
 - Moving media offsite.
 - Backing up the device configuration file, the volume history file, and the dsmserve.opt server options file.
 - **Optional:** Issuing the **PREPARE** command to create the disaster recovery plan file.

Tip:

When you issue the **PREPARE** command, the IBM Storage Protect disaster recovery manager (DRM) function creates one copy of the disaster recovery plan file.

You can manage offsite disaster recovery without using DRM, however, DRM helps to consolidate plans, scripts, and other information that is required during disaster recovery.

Create multiple copies of the plan for safekeeping. For example, keep copies in print, on a USB flash drive, on disk space that is located offsite, or on a remote server. The disaster recovery plan file is moved offsite daily with the tapes. For more information about DRM, see [“Preparing for and recovering from a disaster by using DRM” on page 208](#).

- Configure the following resources at the disaster recovery site:

1. A recovery IBM Storage Protect server. The server at the disaster recovery site must be at the same level as the server on the production site.
2. A tape library to store the media that is shipped from the production site. For more information about offsite recovery locations, see [“Offsite data storage”](#) on page 24.
3. Disk storage space for the database, archive log, active logs, and storage pools.
4. Clients to test restore operations.

About this task

Test the disaster recovery plan and the IBM Storage Protect server recoverability often, in an environment that is similar to the production environment.

Procedure

1. Ensure that tapes are available onsite. Issue the **QUERY LIBVOLUME** command to identify volumes that are checked into an automated library.
2. Back up the database to the onsite tapes by completing the following steps:
 - a. On the **Servers** page of the Operations Center, select the server whose database you want to back up.
 - b. Click **Back Up**, and follow the instructions in the **Back Up Database** window.
3. Copy the following files to the home directory of the server at the recovery site:
 - Disaster recovery plan file
 - Volume history file
 - Device configuration file
 - Optional: dsmserve.opt server options file
4. Move the tape to the offsite recovery location.
5. Restore the server database by using the **DSMSERV RESTORE DB** utility on the recovery server.
6. Issue the **UPDATE VOLUME** command and specify the **ACCESS=DESTROYED** parameter to indicate that an entire volume must be restored.
7. On the recovery server, restore the storage pool volumes by using the **RESTORE STGPOOL** command.

What to do next

Ensure that you can access the data in the library by auditing a tape volume in the restored storage pool to verify that the data is consistent. Issue the **AUDIT VOLUME** command to audit a tape volume. For faster performance, audit restored data only.

Related tasks

[Auditing the volume inventory in a library](#)

You can audit an automated library to ensure that the library volume inventory is consistent with the volumes that are physically in the library. You might want to audit a library if the library volume inventory is distorted due to manual movement of volumes in the library or to database problems.

Related information

[AUDIT VOLUME \(Verify database information for a storage pool volume\)](#)

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

[RESTORE STGPOOL \(Restore storage pool data\)](#)

Restoring the database

If you have the disaster recovery manager (DRM) function enabled and you followed the procedure to prepare for a disaster, you can restore the database after a disaster. If you do not have DRM configured, you can still restore the database, provided that you have the required backup files. If you do not have the required backup files, you can potentially restore the database by recreating 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

You can restore the database to its most current state or to a specified point in time. To recover the database to the time when the database was lost, recover the database to its latest version.

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 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 V7.1 database, an error occurs.

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 todate=04/19/2022
```

Related information

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

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

Numerics

- 3590 tape drive
 - defining device class [20](#)
- 3592 drives and media
 - cleaning [197](#)
 - data encryption [99](#), [122](#)
 - defining device class [20](#)
 - DEVICETYPE parameter [176](#)
 - enabling for WORM media [130](#)
 - mixing drive generations [97](#)

A

- About this publication [vii](#)
- accessibility features [215](#)
- active log capacity [162](#)
- administrative commands
 - AUDIT LIBVOLUME [187](#)
 - CHECKIN LIBVOLUME [176](#), [178](#)
 - CHECKOUT LIBVOLUME [185](#)
 - CLEAN DRIVE [196](#)
 - DEFINE DEVCLASS
 - 3592 [97](#)
 - LTO device classes [94](#)
 - DEFINE DRIVE [92](#)
 - DEFINE LIBRARY [91](#)
 - UPDATE DRIVE [191](#), [192](#)
 - UPDATE LIBVOLUME [185](#)
 - UPDATE VOLUME [181](#)
 - VALIDATE LANFREE [119](#)
- archive log capacity [162](#)
- archive operations
 - scheduling [112](#)
 - specifying rules [109](#)
- AUDIT LIBVOLUME command [187](#)
- auditing
 - library volume inventory [187](#)
- authority level [201](#)
- autochangers [84](#)
- AUTOLABEL parameter for tape volumes [175](#)
- automated library
 - scratch volume [185](#)
- automated library device
 - auditing [187](#)
 - changing volume status [185](#)
 - checking in volumes [176](#)
 - informing server of new volumes [176](#)
 - labeling volumes [174](#)
 - removing volumes [185](#)
 - replacing tape drive [199](#)
 - volume inventory [188](#)

B

- back-end capacity licensing [152](#)
- backup operations

- backup operations (*continued*)
 - modifying scope [112](#)
 - scheduling [112](#)
 - specifying rules [109](#)
- bar code reader [178](#)
- bar-code reader
 - auditing volumes in a library [187](#)
- barcode reader
 - checking in volumes for a library [178](#)
 - labeling volumes in a library [175](#)

C

- capacity, tape [123](#)
- cartridge
 - cleaner cartridge [152](#), [198](#)
 - mixing drive generations [97](#)
- check in
 - cleaner cartridge [198](#)
 - library volume [176](#), [178](#)
 - setting a time interval for volume [126](#)
- CHECKIN LIBVOLUME command [176](#), [177](#)
- CHECKOUT LIBVOLUME command [185](#)
- class, device
 - defining [93](#)
 - FORMAT parameter [124](#)
 - LTO [94](#)
- CLEAN DRIVE command [196](#), [199](#)
- cleaner cartridge
 - checking in [198](#)
 - operations with [152](#), [198](#)
- client acceptor
 - configuring [116](#)
 - restarting [156](#)
 - stopping [156](#)
- client nodes
 - decommissioning [159](#)
 - removing from production [159](#)
- client/server communications
 - configuring [118](#)
- clients
 - adding [107](#)
 - configuring [114](#)
 - configuring to run scheduled operations [116](#)
 - connecting to server [113](#)
 - define schedules [75](#)
 - installing [114](#)
 - managing operations [155](#)
 - protecting [107](#)
 - registering [113](#)
 - selecting software [108](#)
 - upgrading [158](#)
- collocation
 - changing, effect of [170](#)
 - definition [164](#)
 - determining whether to use collocation [164](#)
 - effects on operations [166](#)

- collocation (*continued*)
 - enabling [171](#)
 - enabling for sequential storage pool [164](#)
 - how the server selects volumes when disabled [169](#)
 - planning [171](#)
 - selecting volumes when enabled [167](#)
- collocation of retention storage pools [171](#)
- commands
 - HALT [205](#)
- commands, administrative
 - AUDIT LIBVOLUME [187](#)
 - CHECKIN LIBVOLUME [176](#), [178](#)
 - CHECKOUT LIBVOLUME [185](#)
 - CLEAN DRIVE [196](#)
 - DEFINE DEVCLASS
 - 3592 [97](#)
 - LTO device classes [94](#)
 - DEFINE DRIVE [92](#)
 - DEFINE LIBRARY [91](#)
 - UPDATE DRIVE [191](#), [192](#)
 - UPDATE LIBVOLUME [185](#)
 - UPDATE VOLUME [181](#)
 - VALIDATE LANFREE [119](#)
- configuration
 - changing [156](#)
 - clients [114](#)
- configuring
 - shared library [100](#)
- configuring libraries
 - SCSI [90](#)
- copy retention to tape [67](#), [74](#)

D

- daily checklist of monitoring tasks [135](#)
- data
 - deactivating [161](#)
- data encryption [120](#)
- data migration [2](#)
- data protection with WORM media [129](#)
- data recovery
 - strategy [211](#)
- data retention rules
 - define [57](#)
- database capacity [162](#)
- database restore [212](#)
- deactivation process
 - backup data [161](#)
- decommission process
 - client node [159](#)
- define drive [200](#)
- DEFINE DRIVE command [92](#)
- DEFINE LIBRARY command [91](#)
- DELETE DRIVE [199](#)
- determining
 - the time interval for volume check in [126](#)
- device
 - multiple types in a library [18–20](#)
 - name [78](#)
 - zfcps device driver [87](#)
- device class
 - defining [93](#)
 - FORMAT parameter [124](#)
 - LTO [94](#)

- device diagnostics [129](#)
- device driver
 - configuring [84](#), [86](#), [88](#), [89](#)
 - for automated library devices [76](#)
 - IBM Storage Protect, installing [76](#)
 - installing [75](#)
 - requirements [75](#)
- device drivers
 - installing [79](#)
- device type
 - LTO [94](#)
 - multiple in a single library [18](#), [20](#)
- device, storage
 - device information [129](#)
 - replacing tape drive [199](#)
 - required IBM Storage Protect definitions [20](#)
- devices
 - defining [91](#)
- diagnostics, for device [129](#)
- disability [215](#)
- disaster
 - disaster recovery manager [208](#)
- disaster preparation [208](#)
- disaster recovery [62–64](#), [208](#), [210](#), [211](#)
- disaster recovery manager [62–64](#), [208](#), [210–212](#)
- disaster recovery plan file [208](#)
- DISMOUNT VOLUME command [189](#)
- DLT WORM media [129](#)
- drive
 - cleaning [196](#), [199](#)
 - defining [92](#)
 - detecting changes on a SAN [128](#)
 - element address [92](#)
 - serial number [92](#)
 - updating [191](#), [192](#)
- drive cleaning [196](#)
- DRIVEENCRYPTION parameter
 - 3592 device class [99](#)
 - LTO device class [96](#)
- driver, device
 - configuring [84](#)
 - for automated library devices [76](#)
 - IBM Storage Protect, installing [76](#)
 - installing [75](#)
 - requirements [75](#)
- driver, tape device
 - installing [76](#)
 - requirements [76](#)
- drivers, device [79](#)
- DRM [62–64](#), [208](#), [210–212](#)
- DSMSERV RESTORE DB [212](#)

E

- electronic vaulting [24](#)
- element address [92](#), [180](#)
- email reports
 - configuring [154](#)
- encryption
 - DRIVEENCRYPTION parameter
 - 3592 Generation 2 [99](#)
 - LTO-4 or later [96](#)
 - methods [120](#), [122](#)
 - options [26](#)

- error checking
 - clean drive [199](#)
- error logs
 - evaluating [155](#)

F

- fibre channel devices [83](#)
- fibre channel SAN-attached devices [85](#)
- file name for a device [78](#)
- file systems
 - [preparing, AIX server systems [43](#)
 - planning for [8](#)
 - preparing, Linux server systems [45](#)
 - preparing, Windows server systems [46](#)
- firewall [27](#)
- firewalls
 - configuring communications through [118](#)
- front-end capacity licensing [152](#)

G

- graphical wizard
 - prerequisite RPM files [47](#)

H

- halting
 - server [205](#)
- hardware requirements [3](#)

I

- IBM device drivers
 - configuring [79](#)
 - installing [79](#)
- IBM Documentation [vii](#)
- IBM Storage Protect device driver [76](#)
- IBM Storage Protect device drivers [77](#)
- IBM Storage Protect directories
 - planning for [8](#)
- IBM Storage Protect disaster recovery manager [208](#), [210](#)
- IBM tape device drivers [77](#)
- implementation
 - test operations [133](#)
- installation
 - clients [114](#)
- installing IBM Storage Protect
 - AIX systems [46](#)
 - Linux systems [46](#)
 - Windows systems [48](#)
- installing the operating system
 - AIX server systems [31](#)
 - Linux server systems [32](#)
 - Windows server systems [37](#)
- inventory capacity [162](#)
- issues
 - diagnosing [135](#)

K

- keyboard [215](#)

L

- label
 - automatic labeling in SCSI libraries [175](#)
 - bar code reader [178](#)
 - checking in [178](#)
 - checking media [178](#)
 - overwriting existing labels [174](#), [175](#)
 - sequential storage pools [173](#)
 - volume examples [174](#)
 - volumes using a library device [174](#)
- LABEL LIBVOLUME command
 - identifying drives [174](#)
 - labeling sequential storage pool volumes [174](#)
 - overwriting existing volume labels [174](#)
 - removable media volumes [174](#)
 - using a library device [174](#)
 - volume labeling examples [174](#)
- LAN-free data movement
 - description [16](#), [17](#)
- LDAP
 - password requirements [202](#)
- library
 - adding volumes [176](#)
 - auditing volume inventory [187](#)
 - automated [184](#)
 - configuration [90](#)
 - configure for more than one device type [18](#), [20](#)
 - defining [91](#)
 - detecting changes to, on a SAN [92](#), [128](#)
 - mixing device types [18](#), [20](#), [94](#), [97](#)
 - mode, random or sequential [76](#)
 - SCSI [15](#)
 - serial number [92](#)
 - shared [15](#)
 - sharing among servers [100](#)
 - volume inventory [188](#)
- library client, shared library [16](#), [104](#)
- library manager, shared library [16](#), [103](#)
- library sharing [17](#)
- library storage slot [180](#)
- library storage slots [177](#)
- license compliance
 - verifying [152](#)
- logical block protection
 - enabling [194](#)
 - overview [192](#)
 - read/write operations [195](#)
 - storage pool management [195](#)
 - supported drives [193](#)
- LTO Ultrium devices and media
 - device class, defining and updating [94](#)
 - encryption [96](#), [122](#)
 - WORM [129](#)

M

- maintenance
 - define schedule [58](#)
- maintenance mode
 - start server [204](#)
- maintenance tasks
 - scheduling [163](#)
 - start the server in maintenance mode [206](#)

- managing
 - access levels [204](#)
 - administrators [201](#)
 - authority [201](#)
- managing security [51](#)
- media
 - tape rotation [181](#)
- media incompatibility [151](#)
- media label
 - checking [178](#)
 - for tape [174](#)
 - recording [174](#)
- messages
 - for automated libraries [189](#)
- migrating drives [201](#)
- mixed device types in a library [18–20](#), [94](#), [97](#)
- mode
 - library (random or sequential) [76](#)
- monitoring
 - daily checklist [135](#)
 - goals [135](#)
 - periodic checklist [145](#)
 - tasks
 - daily checklist [135](#)
 - periodic checklist [145](#)
- mount
 - library [125](#)
 - limit [125](#)
 - operations [189](#)
 - query [189](#)
 - retention period [126](#)
 - wait period [126](#)
- mount point
 - preemption [127](#)
 - relationship to mount limit in a device class [125](#)
- move data [62](#)
- move media [62–64](#)
- MOVE RETMEDIA [69](#), [71](#)
- multipath I/O
 - configure for AIX systems [38](#)
 - configure for Linux systems [39](#)
 - configure for Windows systems [40](#)

N

- name of device [78](#)
- network bandwidth [2](#)
- new tape drive [199](#)
- NOPREEMPT server option [127](#)

O

- offsite storage [24](#)
- offsite vaulting [24](#)
- offsite volume [63](#), [69](#)
- onsite volume [64](#), [71](#)
- operating system
 - install on AIX server systems [31](#)
 - install on Linux server systems [32](#)
 - install on Windows server systems [37](#)
 - security [204](#)
- Operations Center
 - configure [54](#)

- Operations Center (*continued*)
 - secure communications [55](#)
- option, server
 - NOPREEMPT [127](#)
- options
 - set for server [50](#)
- outage
 - prepare [207](#)

P

- passthru driver [77](#)
- password requirements
 - LDAP [202](#)
- passwords
 - changing [202](#)
 - resetting [157](#)
- paths
 - defining [91](#)
- performance
 - volume frequently used, improve with longer mount retention [126](#)
- periodic checklist of monitoring tasks [145](#)
- planning solutions
 - tape [1](#)
- planning worksheet [8](#)
- policies
 - editing [111](#)
 - specifying [109](#)
 - viewing [110](#)
- policy domains
 - specifying [109](#)
- pool, storage
 - 3592, special considerations for [97](#)
 - determining whether to use collocation [164](#)
 - LTO Ultrium, special considerations for [94](#)
- preemption
 - mount point [127](#)
 - volume access [128](#)
- privilege class
 - system privilege [201](#)
- processor value unit (PVU) licensing [152](#)
- product license
 - register [57](#)
- protecting your data [129](#)
- publications [vii](#)

Q

- QUERY SAN [129](#)

R

- random mode for libraries [76](#)
- reconfiguration tasks
 - start the server in maintenance mode [206](#)
- recovery drill [211](#)
- registration
 - clients [113](#)
- remove drive [199](#)
- replace drive [200](#)
- replacing tape drive [199](#)
- reports

- reports (*continued*)
 - email
 - configuring [154](#)
- restricting
 - user access [204](#)
- retention volumes [67](#), [69](#), [71](#), [74](#)
- RPM files
 - install for graphical wizard [47](#)
- rules
 - editing [111](#)
 - specifying
 - backup and archive operations [109](#)
 - viewing [110](#)

S

- SAN (storage area network)
 - client access to devices [16](#), [17](#)
 - device changes, detecting [128](#)
 - LAN-free data movement [16](#), [17](#)
 - sharing a library among servers [16](#), [100](#)
 - storage agent role [16](#), [17](#)
- scale capacity [99](#)
- scheduled activities
 - tuning [163](#)
- schedules
 - backup and archive operations [112](#)
- scratch volume [185](#)
- scratch volumes [183](#)
- SCSI
 - automatic labeling of volumes [175](#)
 - library with different tape technologies [97](#)
- SCSI devices [83](#)
- SCSI libraries
 - define a library client [101](#), [102](#)
 - define a library server [101](#), [102](#)
- secure communications
 - configure with SSL and TLS [54](#)
- security
 - data encryption
 - 3592 Generation 2 [99](#)
 - IBM LTO Generation 4 [122](#)
 - IBM LTO Generation 4 or later [96](#)
 - Oracle StorageTek T10000B [122](#)
 - Oracle StorageTek T10000C [122](#)
 - Oracle StorageTek T10000D [122](#)
 - data encryption, 3592 Generation 2, TS1120, TS1130, TS1140, TS1150 [122](#)
- sequential mode for libraries [76](#)
- serial number
 - automatic detection by the server [92](#), [128](#)
 - for a drive [92](#)
 - for a library [92](#)
- server
 - configure [49](#)
 - create user ID for [41](#)
 - define maintenance schedule [58](#)
 - plan upgrade [206](#)
 - set options [50](#)
 - start in maintenance mode [204](#)
 - stop [205](#)
- server option
 - NOPREEMPT [127](#)
- servers

- servers (*continued*)
 - start in maintenance mode [206](#)
- setting
 - library mode [76](#)
 - time interval for checking in volumes [126](#)
- shared SCSI library [100](#)
- shutting down
 - server [205](#)
- software
 - selecting [108](#)
- software requirements [6](#)
- solution
 - expanding [107](#)
- Sony WORM media (AIT50 and AIT100) [129](#)
- special file names [78](#)
- SSL [54](#)
- starting server
 - maintenance mode [204](#)
- status reports
 - obtaining [154](#)
- stopping
 - server [205](#)
- storage
 - planning for [13](#), [15](#)
- storage agent [16](#), [17](#)
- storage area network (SAN)
 - client access to devices [16](#), [17](#)
 - device changes, detecting [128](#)
 - LAN-free data movement [16](#), [17](#)
 - sharing a library among servers [16](#), [100](#)
 - storage agent role [16](#), [17](#)
- storage configuration
 - planning for [8](#)
- storage devices [93](#)
- storage hardware
 - configure [30](#)
- storage pool
 - 3592, special considerations for [97](#)
 - determining whether to use collocation [164](#)
 - LTO Ultrium, special considerations for [94](#)
- storage pool capacity [2](#)
- storage pool hierarchies
 - planning [21](#)
 - setting up [105](#)
- storage space
 - releasing [161](#)
- storage volume
 - labeling sequential access [173](#)
 - preparing sequential access [173](#)
- swapping volumes in automated library [179](#)
- system requirements
 - hardware [3](#)
- system status
 - tracking [154](#)
- system update
 - prepare [207](#)

T

- tape
 - capacity [123](#)
 - compatibility between drives [199](#)
 - recording format [124](#)
 - rotation [181](#)

- tape (*continued*)
 - setting mount retention period [126](#)
- tape device driver
 - installing [76](#)
 - requirements [76](#)
- tape drive, replacing [199](#)
- tape drives [2](#)
- tape labels
 - overwriting [87](#)
- tape requirements [2](#)
- tape solution
 - planning for [1](#)
- time interval, setting for checking in volumes [126](#)
- TLS [54](#)
- troubleshooting
 - administrator IDs [157](#)
 - errors in client operations [155](#)
 - locked client nodes [157](#)
 - password issues [157](#)
- tsmdlst utility [129](#)
- type, device
 - LTO [94](#)
 - multiple in a single library [18](#), [20](#)

U

- Ultrium, LTO device type
 - device class, defining and updating [94](#)
 - encryption [96](#), [122](#)
 - WORM [129](#)
- unavailable access mode
 - marked with PERMANENT parameter [189](#)
- UPDATE DRIVE command [191](#), [192](#)
- UPDATE LIBVOLUME command [185](#)
- upgrade
 - server [206](#)
- upgrading tape drives [199](#)
- user ID
 - create for server [41](#)

V

- VALIDATE LANFREE command [119](#)
- validating data
 - logical block protection [192](#)
- volume capacity [124](#)
- volumes
 - access preemption [128](#)
 - access, controlling [181](#)
 - auditing [187](#)
 - automated library inventory [188](#)
 - checking in new volumes to library [176](#)
 - checking out [185](#)
 - determining which are mounted [189](#)
 - dismounting [189](#)
 - inventory maintenance [181](#)
 - managing [184](#)
 - mount retention time [126](#)
 - removing from a library [185](#)
 - sequential storage pools [173](#)
 - swapping [179](#)
 - updating [185](#)

W

- WORM devices and media
 - DLT WORM [129](#)
 - IBM 3592 [129](#)
 - LTO WORM [129](#)
 - maintaining volumes in a library [183](#)
 - Oracle StorageTek T10000B drives [130](#)
 - Oracle StorageTek T10000C drives [130](#)
 - Oracle StorageTek T10000D drives [130](#)
 - Quantum LTO3 [129](#)
 - Sony AIT50 and AIT100 [129](#)
 - special considerations for WORM media [129](#)
 - VolSafe
 - considerations for media [129](#)
- worm volumes [131](#)



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