

IBM Tivoli Storage Manager
Version 7.1.7

*Introduction to Data Protection
Solutions*



IBM Tivoli Storage Manager
Version 7.1.7

*Introduction to Data Protection
Solutions*



Note:

Before you use this information and the product it supports, read the information in “Notices” on page 53.

Fifth edition (September 2016)

This edition applies to version 7, release 1, modification 7 of IBM Tivoli Storage Manager (product numbers 5608-E01, 5608-E02, 5608-E03), and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1993, 2016.

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

Contents

About this publication	v
Who should read this guide	v
Publications	v

Part 1. Tivoli Storage Manager concepts **1**

Chapter 1. Tivoli Storage Manager overview	3
Data protection components	3
Data protection services	4
Processes for managing data protection with Tivoli Storage Manager	7
User interfaces for the Tivoli Storage Manager environment	10

Chapter 2. Data storage concepts in Tivoli Storage Manager.	13
Types of storage devices	13
Data storage in storage pools	17
Data transport to storage across networks	22

Chapter 3. Data protection strategies with Tivoli Storage Manager	27
Strategies to minimize the use of storage space for backups	27
Strategies for disaster protection	28
Strategies for disaster recovery with Tivoli Storage Manager	31

Part 2. Tivoli Storage Manager solutions for data protection **35**

Chapter 4. Disk-based implementation of a data protection solution for a single site	37
---	-----------

Chapter 5. Disk-based implementation of a data protection solution for multiple sites	39
--	-----------

Chapter 6. Tape-based implementation of a data protection solution	41
---	-----------

Chapter 7. Appliance-based implementation of a data protection solution for multiple sites	43
---	-----------

Chapter 8. Comparison of data protection solutions	45
---	-----------

Chapter 9. Roadmap for implementing a data protection solution	47
---	-----------

Part 3. Appendixes. **49**

Appendix. Accessibility features for the Tivoli Storage Manager product family	51
---	-----------

Notices	53
--------------------------	-----------

Glossary	57
---------------------------	-----------

Index	59
------------------------	-----------

About this publication

This publication provides an overview of IBM® Tivoli® Storage Manager concepts and data protection solutions that use best practices for Tivoli Storage Manager. A feature comparison chart helps you select the best solution for your organization's needs.

Who should read this guide

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

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

Publications

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

To view IBM product documentation, see IBM Knowledge Center.

Part 1. Tivoli Storage Manager concepts

IBM Tivoli Storage Manager provides a comprehensive data protection environment.

Chapter 1. Tivoli Storage Manager overview

IBM Tivoli Storage Manager provides centralized, automated data protection that helps to reduce data loss and manage compliance with data retention and availability requirements.

Data protection components

The data protection solutions that Tivoli Storage Manager provides consist of a server, client systems and applications, and storage media. Tivoli Storage Manager provides management interfaces for monitoring and reporting the data protection status.

Server

Client systems send data to the server to be stored as backups or archived data. The server includes an *inventory*, which is a repository of information about client data.

The inventory includes the following components:

Database

Information about each file, logical volume, or database that the server backs up, archives, or migrates is stored in the server database. The server database also contains information about the policy and schedules for data protection services.

Recovery log

Records of database transactions are kept in this log. The database uses the recovery log to ensure data consistency in the database.

Client systems and applications

Clients are applications, virtual machines, and systems that must be protected. The clients send data to the server, as shown in Figure 1.




Figure 1. Components in the data protection solution

Client software

For Tivoli Storage Manager to protect client data, the appropriate software must be installed on the client system and the client must be registered with the server.

Client nodes

A *client node* is equivalent to a computer, virtual machine, or application, such as a backup-archive client that is installed on a workstation for file system backups. Each client node must be registered with the server. Multiple nodes can be registered on a single computer.

Storage media

The server stores client data to storage media. The following types of media are used:

Storage devices

The server can write data to hard disk drives, disk arrays and subsystems, stand-alone tape drives, tape libraries, and other types of random-access and sequential-access storage. Storage devices can be connected directly to the server or connected through a local area network (LAN) or a storage area network (SAN).

Storage pools

Storage devices that are connected to the server are grouped into *storage pools*. Each storage pool represents a set of storage devices of the same media type, such as disk or tape drives. Tivoli Storage Manager stores all of the client data in storage pools. You can organize storage pools into a *hierarchy*, so that data storage can transfer from disk storage to lower-cost storage such as tape devices.

Data protection services

Tivoli Storage Manager provides data protection services to store and recover data from various types of clients. The data protection services are implemented through policies that are defined on the server. You can use client scheduling to automate the data protection services.

Types of data protection services

Tivoli Storage Manager provides services to store and recover client data as shown in Figure 2 on page 5.




Figure 2. Data protection services

Tivoli Storage Manager provides the following types of data protection services:

Back up and restore services

You run a backup process to create a copy of a *data object* that can be used for recovery if the original data object is lost. A data object can be a file, a directory, or a user-defined data object, such as a database.

To minimize the use of system resources during the backup operation, Tivoli Storage Manager uses the *progressive incremental backup* method. For this backup method, a first full backup of all data objects is created and in subsequent backup operations only changed data is moved to storage. Compared to incremental and differential backup methods that require taking periodic full backups, the progressive incremental backup method provides the following benefits:

- Reduces data redundancy
- Uses less network bandwidth
- Requires less storage pool space

To further reduce storage capacity requirements and network bandwidth usage, Tivoli Storage Manager includes *data deduplication* for data backups. The data deduplication technique removes duplicate data extents from backups.

You run a restore process to copy an object from a storage pool to the client. You can restore a single file, all files in a directory, or all of the data on a computer.

Archive and retrieve services

You use the archive service to preserve data that must be stored for a long time, such as for regulatory compliance. The archive service provides the following features:

- When you archive data, you specify how long the data must be stored.
- You can request that files and directories are copied to long-term storage on media. For example, you might choose to store this data on a tape device, which can reduce the cost of storage.
- You can specify that the original files are erased from the client after the files are archived.

The retrieve service provides the following features:

- When you retrieve data, the data is copied from a storage pool to a client node.
- The retrieve operation does not affect the archive copy in the storage pool.

Migrate and recall services

You use migrate and recall services to manage space on client systems. The goal of space management is to maximize available media capacity for new data and to minimize access time to data. You can migrate data to server storage to maintain sufficient free storage space on a local file system. You can store migrated data in the following ways:

- On disk storage for long-term storage
- In a *virtual tape library* (VTL) for fast recall of files

You can recall files to the client node on demand, either automatically or selectively.

Types of client data that can be protected

You can protect data for the following types of clients with Tivoli Storage Manager:

Application clients

Tivoli Storage Manager can protect data for specific products or applications. These clients are called *application clients*. To protect the *structured data* for these clients, in other words the data in database fields, you must back up components that are specific to the application. Tivoli Storage Manager can protect the following applications:

- IBM Tivoli Storage Manager for Enterprise Resource Planning clients:
 - Data protection for SAP HANA
 - Data protection for SAP for DB2®
 - Data protection for SAP for Oracle
- IBM Tivoli Storage Manager for Databases clients:
 - Data protection for Microsoft SQL server
 - Data protection for Oracle
- IBM Tivoli Storage Manager for Mail clients:
 - Data protection for IBM Domino®
 - Data protection for Microsoft Exchange Server

Virtual machines

Virtual machines that are backed up by using application client software that is installed on the virtual machine. In the Tivoli Storage Manager environment, a virtual machine can be protected by the IBM Tivoli Storage Manager for Virtual Environments.

System clients

The following Tivoli Storage Manager clients are called *system clients*:

- All clients that back up data in files and directories, in other words *unstructured data*, such as backup-archive clients and API clients that are installed on workstations.
- A server that is included in a server-to-server virtual volume configuration.
- A virtual machine that is backed up by using backup-archive client software that is installed on the virtual machine.

Processes for managing data protection with Tivoli Storage Manager

The Tivoli Storage Manager server inventory has a key role in the processes for data protection. You define policies that the server uses to manage data storage.

Data management process

Figure 3 shows the Tivoli Storage Manager data management process.

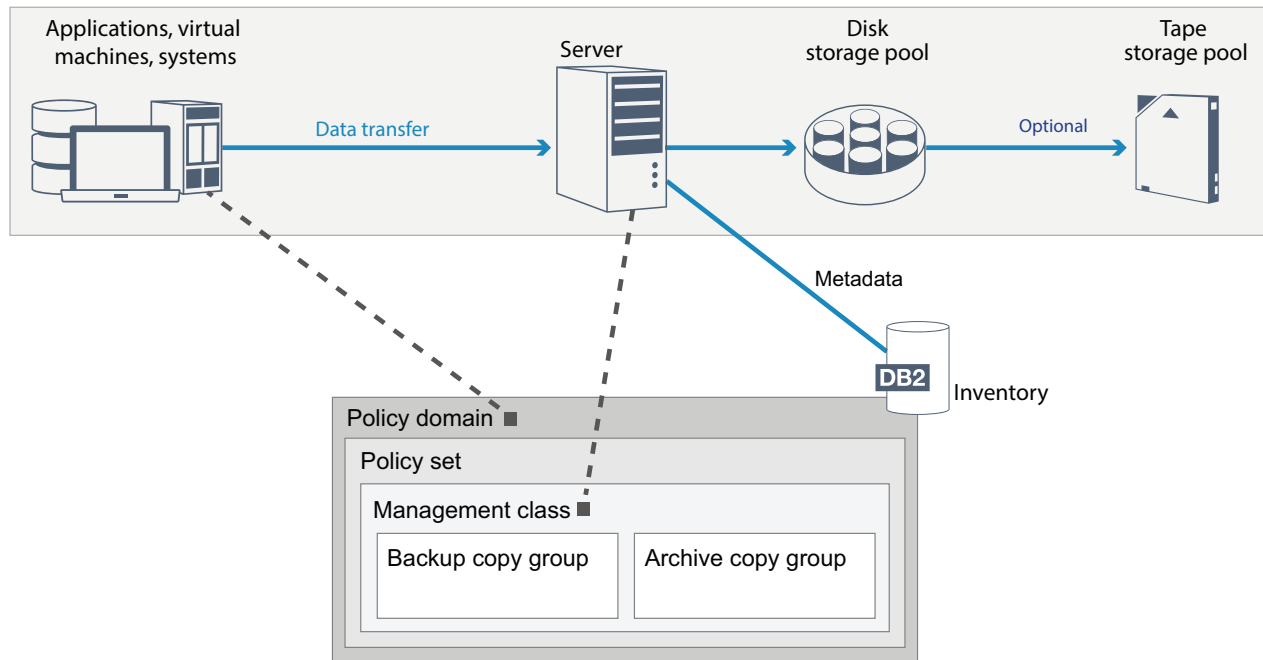


Figure 3. Data management process

Tivoli Storage Manager uses policies to control how the server stores and manages data objects on various types of storage devices and media. You associate a client with a policy domain that contains one active policy set. When a client backs up, archives, or migrates a file, the file is bound to a management class in the active policy set of the policy domain. The management class and the backup and archive copy groups specify where files are stored and how they are managed. If you set up server storage in a hierarchy, you can migrate files to different storage pools.

Inventory components

The following inventory components are key to the operation of the server:

Server database

The server database contains information about client data and server operations. The database stores information about client data, called *metadata*. Information about client data includes the file name, file size, file owner, management class, copy group, and location of the file in server storage. The database includes the following information that is necessary for the operation of the server:

- Definitions of client nodes and administrators
- Policies and schedules
- Server settings

- Records of server operations, such as activity logs and event records
- Intermediate results for administrative queries

Recovery log

The server records database transactions in the recovery log. The recovery log helps to ensure that a failure does not leave the database in an inconsistent state. The recovery log is also used to maintain consistency across server start operations. The recovery log consists of the following logs:

Active log

This log records current transactions on the server. This information is required to start the server and database after a disaster.

Log mirror (optional)

The active log mirror is a copy of the active log that can be used if the active log files cannot be read. All changes that are made to the active log are also written to a log mirror. You can set up one active log mirror.

Archive log

The archive log contains copies of closed log files that were in the active log. The archive log is included in database backups and is used for recovery of the server database. Archive log files that are included in a database backup are automatically pruned after a full database backup cycle is complete. The archive log must have enough space to store the log files for database backups.

Archive failover log (optional)

The archive failover log, also called a secondary archive log, is the directory that the server uses to store archive log files when the archive log directory is full.

Policy-based data management

In the Tivoli Storage Manager environment, a *policy* for data protection management contains rules that determine how client data is stored and managed. The primary purpose of a policy is to implement the following data management objectives:

- Control which storage pool client data is initially stored in
- Define retention criteria that controls how many copies of objects are stored
- Define how long copies of objects are retained

Policy-based data management helps you to focus on the business requirements for protecting data rather than on managing storage devices and media. Administrators define policies and assign client nodes to a *policy domain*.

Depending on your business needs, you can have one policy or many. In a business organization, for example, different departments with different types of data can have customized storage management plans. Policies can be updated, and the updates can be applied to data that is already managed.

When you install Tivoli Storage Manager, a default policy that is named STANDARD is already defined. The STANDARD policy provides basic backup protection for user workstations. To provide different levels of service for different clients, you can add to the default policy or create a new policy.

You create policies by defining the following policy components:

Policy domain

The policy domain is the primary organizational method of grouping client nodes that share common rules for data management. Although a client node can be defined to more than one server, the client node can be defined to only one policy domain on each server.

Policy set

A *policy set* is a number of policies that are grouped so that the policy for the client nodes in the domain can be activated or deactivated as required. An administrator uses a policy set to implement different management classes based on business and user needs. A policy domain can contain multiple policy sets, but only one policy set can be active in the domain. Each policy set contains a default management class and any number of extra management classes.

Management class

A *management class* is a policy object that you can bind to each category of data to specify how the server manages the data. There can be one or more management classes. One management class is assigned to be the default management class that is used by clients unless they specifically override the default to use a specific management class.

The management class can contain a backup copy group, an archive copy group, and space management attributes. A copy group determines how the server manages backup versions or archived copies of the file. The space management attributes determine whether the file is eligible for migration by the space manager client to server storage, and under what conditions the file is migrated.

Copy group

A *copy group* is a set of attributes in a management class that controls the following factors:

- Where the server stores versions of backed up files or archive copies
- How long the server keeps versions of backed up files or archive copies
- How many versions of backup copies are retained
- What method to use to generate versions of backed up files or archive copies

Security management

Tivoli Storage Manager includes security features for registration of administrators and users. After administrators are registered, they must be granted authority by being assigned one or more administrative privilege classes. An administrator with system privilege can perform any server function. Administrators with policy, storage, operator, or node privileges can perform subsets of server functions. The server can be accessed by using the following methods, each controlled with a password:

- Administrator access to manage the server
- Client access to nodes to store and retrieve data

Also included are features that can help to ensure security when clients connect to the server. Depending on business requirements, as an administrator, you can choose one of the following client registration methods:

Open registration

When the client first connects to the server, the user is requested for a node name, password, and contact information. Open registration provides the user with following default settings:

- The client node is assigned to the STANDARD policy domain.
- The user can define whether files are compressed to decrease the amount of data that is sent over networks and the space that is occupied by the data in storage.
- The user can delete archived copies of files from server storage, but not backup versions of files.

Closed registration

Closed registration is the default method for client registration to the server. For this type of registration, an administrator registers all clients. The administrator can implement the following settings:

- Assign the node to any policy domain
- Determine whether the user can use compression or not, or if the user can choose
- Control whether the user can delete backed up files or archived files

You can add more protection for your data and passwords by using Secure Sockets Layer (SSL). SSL is the standard technology that you use to create encrypted sessions for servers and clients, and provides a secure channel to communicate over open communication paths. With SSL, the identity of the server is verified by using digital certificates. If you authenticate with a Lightweight Directory Access Protocol (LDAP) server, passwords between the server and the LDAP server are protected by Transport Layer Security (TLS). The TLS protocol is the successor to the SSL protocol. When a server and client communicate, TLS ensures that third parties cannot intercept messages.

User interfaces for the Tivoli Storage Manager environment

For monitoring and configuration tasks, Tivoli Storage Manager provides various interfaces, including the Operations Center, a command-line interface, and an SQL administrative interface.

Interfaces for data storage management

The Operations Center is the primary interface for administrators to monitor and administer servers. A key benefit of the Operations Center is that you can monitor multiple servers, as shown in Figure 4 on page 11. You can also monitor and administer Tivoli Storage Manager from a command-line administrative interface.




Figure 4. User interfaces for data storage management

You use the following interfaces to interact with Tivoli Storage Manager:

Operations Center

The Operations Center provides web and mobile access to status information about the Tivoli Storage Manager environment. You can use the Operations Center to complete monitoring and certain administration tasks, for example:

- You can monitor multiple servers and clients.
- You can monitor the transaction activity for specific components in the data path, such as the server database, the recovery log, storage devices, and storage pools.

Command-line interface

You can use a command-line interface to run administration tasks for servers. You can access the command-line interface through either the Tivoli Storage Manager administrative client or the Operations Center.

Access to information in the server database by using SQL statements

You can use SQL SELECT statements to query the server database and display the results. Third-party SQL tools are available to aid administrators in database management.

Interfaces for client activity management

Tivoli Storage Manager provides the following types of interfaces for managing client activity:

- An application programming interface (API)
- Graphical user interfaces for clients
- Browser interface for the backup-archive client
- Command-line interfaces for clients

Chapter 2. Data storage concepts in Tivoli Storage Manager

Tivoli Storage Manager provides functions to store data in a range of device and media storage.

To make storage devices available to the server, you must attach the storage devices and map storage pools to device classes, libraries, and drives.

Types of storage devices

You can use various storage devices with Tivoli Storage Manager to meet specific data protection goals.

Storage devices and storage objects

The Tivoli Storage Manager server can connect to a combination of manual and automated storage devices. You can connect the following types of storage devices to Tivoli Storage Manager:

- Disk devices that are directly attached, SAN-attached, or network attached
- Physical tape devices that are either manually operated or automated
- Virtual tape devices
- Cloud object storage

Tivoli Storage Manager represents physical storage devices and media with storage objects that you define in the server database. Storage objects classify available storage resources and manage migration from one storage pool to another. Table 1 describes the storage objects in the server storage environment.

Table 1. Storage objects and representations

Storage object	What the object represents
Volume	A discrete unit of storage on disk, tape, or other storage media. Each volume is associated with a single storage pool.
Storage pool	A set of storage volumes or containers that is the destination that is used to store client data. Tivoli Storage Manager uses the following types of storage pool: <ul style="list-style-type: none">• Directory-container storage pools• Cloud-container storage pools• Sequential-access storage pools that are associated with a device class• Random-access storage pools that are associated with a device class
Container	A data storage location, for example, a file, directory, or device.
Container storage pool	A primary storage pool that a server uses to store data. Data is stored in containers in file system directories or in cloud storage. Data is deduplicated, if necessary, as the server writes data to the storage pool.

Table 1. Storage objects and representations (continued)

Storage object	What the object represents
Device class	The type of storage device that can use the volumes that are defined in a sequential-access or random-access storage pool. Each device class of removable media type is associated with a single library.
Library	A storage device. For example, a library can represent a stand-alone drive, a set of stand-alone drives, a multiple-drive automated device, or a set of drives that is controlled by a media manager.
Drive	An object of a tape library device that provides the capability to read and write data to tape library media. Each drive is associated with a single library.
Path	The specification of the data source and the device destination. Before a storage device can be used, a path must be defined between the device and the source server that is moving data.
Data mover	A SAN-attached device that is used to transfer client data. A data mover is used only in a data transfer where the server is not present, such as in a Network Data Management Protocol (NDMP) environment. Data movers transfer data between storage devices without using significant server, client, or network resources.
Server	A server that is managed by another Tivoli Storage Manager server.

The administrator defines the storage objects in the logical layer of the server, as illustrated in Figure 5 on page 15.




Figure 5. Storage objects

Disk devices

You can store client data on disk devices with the following types of volume:

- Directories in directory-container storage pools
- Random-access volumes of device type DISK
- Sequential-access volumes of device type FILE

Tivoli Storage Manager offers the following features when you use directory-container storage pools for data storage:

- You can apply data deduplication and disk caching techniques to maximize data storage usage.
- You can retrieve data from disk much faster than you can retrieve data from tape storage.

Physical tape devices

In a physical tape library, the storage capacity is defined in terms of the total number of volumes in the library. Physical tape devices can be used for the following activities:

- Storing client data that is backed up, archived, or migrated from client nodes
- Storing database backups
- Exporting data to another server or offsite storage

Moving data to tape provides the following benefits:

- You can keep data for clients on a disk device at the same time that the data is moved to tape.
- You can improve tape drive performance by streaming the data migration from disk to tape.
- You can spread out the times when the drives are in use to improve the efficiency of the tape drives.
- You can move data on tape to off-site vaults.
- You can limit power consumption because tape devices do not consume power after data is written to tape.
- You can apply encryption that is provided by the tape drive hardware to protect the data on tape.

Compared to equivalent disk and virtual tape storage, the unit cost to store data tends to be much less for physical tape devices.

Virtual tape libraries

A virtual tape library (VTL) does not use physical tape media. When you use VTL storage, you emulate the access mechanisms of tape hardware. In a VTL, you can define volumes and drives to provide greater flexibility for the storage environment. The storage capacity of a VTL is defined in terms of total available disk space. You can increase or decrease the number and size of volumes on disk.

Defining a VTL to the Tivoli Storage Manager server can improve performance because the server handles mount point processing for VTLs differently than for real tape libraries. Although the logical limitations of tape devices are still present, the physical limitations for tape hardware are not applicable to a VTL thus affording better scalability. You can use the Tivoli Storage Manager VTL when the following conditions are met:

- Only one type and generation of drive and media is emulated in the VTL.
- Every server and storage agent with access to the VTL has paths that are defined for all drives in the library.

Data storage in storage pools

Logical storage pools are the principal components in the Tivoli Storage Manager model of data storage. You can optimize the usage of storage devices by manipulating the properties of storage pools and volumes.

Types of storage pools

The group of storage pools that you set up for the server is called *server storage*. You can define the following types of storage pools in server storage:

Primary storage pools

A named set of volumes that the server uses to store backup versions of files, archive copies of files, and files that are migrated from client nodes.

Copy storage pools

A named set of volumes that contain copies of files that reside in primary storage pools. Copy storage pools are used only to back up the data that is stored in primary storage pools. A copy storage pool cannot be a destination for a backup copy group, an archive copy group, or a management class for space-managed files.

Container-copy storage pools

A named set of volumes that contain a copy of data extents that reside in directory-container storage pools. Container-copy storage pools are used only to protect the data that is stored in directory-container storage pools.

Active-data storage pools

A named set of storage pool volumes that contain only active versions of client backup data.

Primary storage pools

When you restore, retrieve, recall, or export file data, the requested file is obtained from a primary storage pool. Depending on the type of primary storage pool, the storage pools can be onsite or offsite. You can arrange primary storage pools in a storage hierarchy so that data can be transferred from disk storage to lower-cost storage such as tape devices. Figure 6 on page 18 illustrates the concept of primary storage pools.




Figure 6. Primary storage pools

You can define the following types of primary storage pool:

Directory-container storage pools

A storage pool that the server uses to store data in containers in storage pool directories. Data that is stored in a directory-container storage pool can use either inline data deduplication, client-side data deduplication, inline compression, or client-side compression. Inline data deduplication or inline compression reduces data at the time it is stored. By using directory-container storage pools, you remove the need for volume reclamation, which improves server performance and reduces the cost of storage hardware. You can protect and repair data in directory-container storage pools at the level of the storage pool.

Restriction: You cannot use any of the following functions with directory-container storage pools:

- Migration
- Reclamation
- Aggregation

- Collocation
- Simultaneous-write
- Storage pool backup
- Virtual volumes

Cloud-container storage pools

A storage pool that a server uses to store data in cloud storage. The cloud storage can be on premises or off premises. The cloud-container storage pools that are provided by Tivoli Storage Manager can store data to cloud storage that is object-based. By storing data in cloud-container storage pools, you can exploit the cost per unit advantages that clouds offer along with the scaling capabilities that cloud storage provides. Tivoli Storage Manager manages the credentials, security, read and write I/Os, and the lifecycle for data that is stored to the cloud. When cloud-container storage pools are implemented on the server, you can write directly to the cloud by configuring a cloud-container storage pool with the cloud credentials. Data that is stored in a cloud-container storage pool can use both inline data deduplication and inline compression. The server writes deduplicated and encrypted data directly to the cloud. You can back up and restore data or archive and retrieve data directly from the cloud-container storage pool.

You can define the following types of cloud-container storage pools:

On premises

You can use the on premises type of cloud-container storage pool to store data in a private cloud, for more security and maximum control over your data. The disadvantages of a private cloud are higher costs due to hardware requirements and onsite maintenance.

Off premises

You can use the off premises type of cloud-container storage pool to store data in a public cloud. The advantage of using a public cloud is that you can achieve lower costs than for a private cloud, for example by eliminating maintenance. However, you must balance this benefit against possible performance issues due to connection speeds and reduced control over your data.

Storage pools that are associated with device classes

You can define a primary storage pool to use the following types of storage devices:

DISK device class

In a DISK device type of storage pool, data is stored in random access disk blocks. You can use caching in DISK storage pools to increase client restore performance with some limitations on server processing. Space allocation and tracking by blocks uses more database storage space and requires more processing power than allocation and tracking by volume.

FILE device class

In a FILE device type of storage pool, files are stored in sequential volumes for better sequential performance than for storage in disk blocks. To the server, these files have the characteristics of a tape volume so that this type of storage pool is better suited for migration to tape. FILE volumes are useful for *electronic vaulting*, where data is transferred electronically to a remote site rather than

by physical shipment of tape. In general, this type of storage pool is preferred over DISK storage pools.

The server uses the following default random-access primary storage pools:

ARCHIVEPOOL

In the STANDARD policy, this storage pool is the destination for files that are archived from client nodes.

BACKUPPOOL

In the STANDARD policy, this storage pool is the destination for files that are backed up from client nodes.

SPACEMGPOOL

This storage pool is for space-managed files that are migrated from Tivoli Storage Manager for Space Management client nodes.

Copy storage pools

Copy storage pools contain active and inactive versions of data that is backed up from primary storage pools. A directory-container storage pool cannot be used as a copy storage pool. In addition, data from a directory-container storage pool cannot be copied into a copy storage pool. To protect directory-container storage pools, copy the data to a container-copy storage pool. Figure 7 illustrates the concept of copy storage pools.

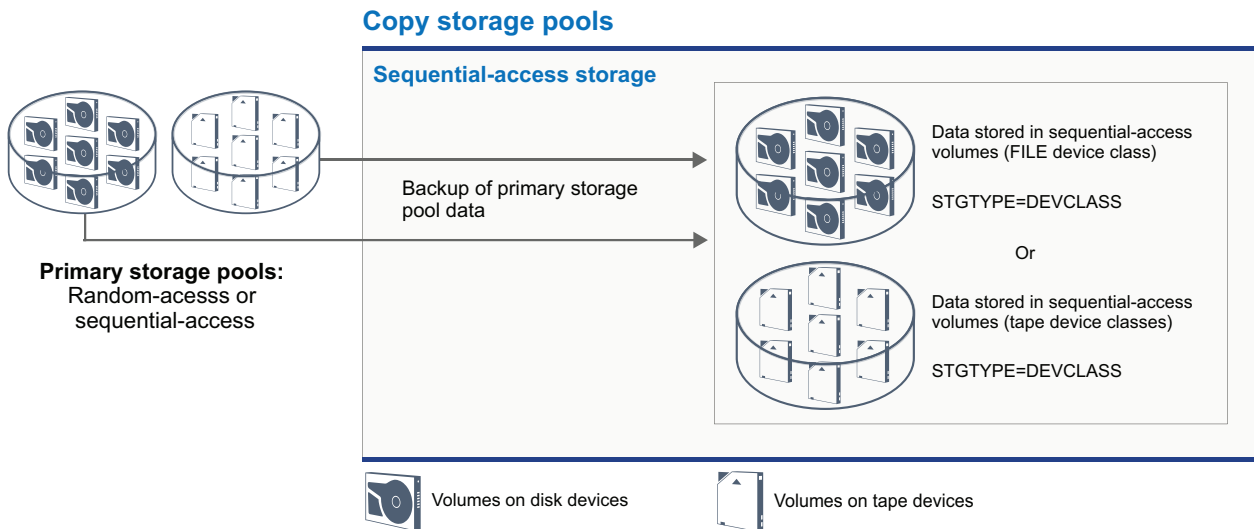


Figure 7. Copy storage pools

Copy storage pools provide a means of recovering from disasters or media failures. For example, when a client attempts to retrieve a damaged file from the primary storage pool, the client can restore the data from the copy storage pool.

You can move the volumes of copy storage pools offsite and still have the server track the volumes. Moving these volumes offsite provides a means of recovering from an onsite disaster. A copy storage pool can use sequential-access storage only, such as a tape device class or FILE device class.

Container-copy storage pools

A server can protect a directory-container storage pool by storing copies of the data in a container-copy storage pool. Data in container-copy storage pools is stored on tape volumes. Damaged data in directory-container storage pools can be repaired by using deduplicated extents in container-copy storage pools.

Container-copy storage pools can be used to repair minor to moderate directory-container storage pool damage, which includes damaged containers or directories. However, replication is the only way to provide complete disaster recovery protection for directory-container storage pools. With replication, you can directly restore client data from the target server if the source server is unavailable.

Figure 8 illustrates the concept of container-copy storage pools.

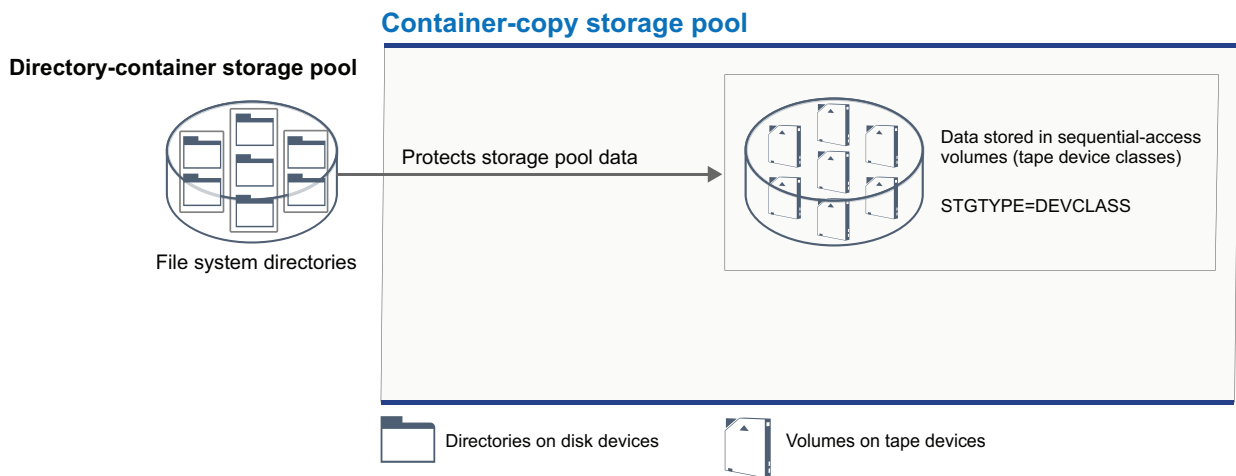


Figure 8. Container-copy storage pools

Depending on the resources and requirements of your site, the ability to copy directory-container storage pools to tape has the following benefits:

- You avoid maintaining another server and more disk storage space.
- Data is copied to storage pools that are defined on the server. Performance is not dependent on, or affected by, the network connection between servers.

Active-data storage pools

An active-data pool contains only active versions of client backup data. In this case, the server does not have to position past inactive files that do not have to be restored. A directory-container storage pool cannot be used as an active-data storage pool. You use active-data pools to improve the efficiency of data storage and restore operations, for example this type of storage pool can help you to achieve the following objectives:

- Increase the speed of client data restore operations
- Reduce the number of onsite or offsite storage volumes
- Reduce the amount of data that is transferred when you copy or restore files that are vaulted electronically in a remote location

Data that is migrated by hierarchical storage management (HSM) clients and archive data are not permitted in active-data pools. As updated versions of backup

data are stored in active-data pools, older versions are removed as the remaining data is consolidated from many sequential-access volumes onto fewer, new sequential-access volumes. Figure 9 illustrates the concept of active-data storage pools.

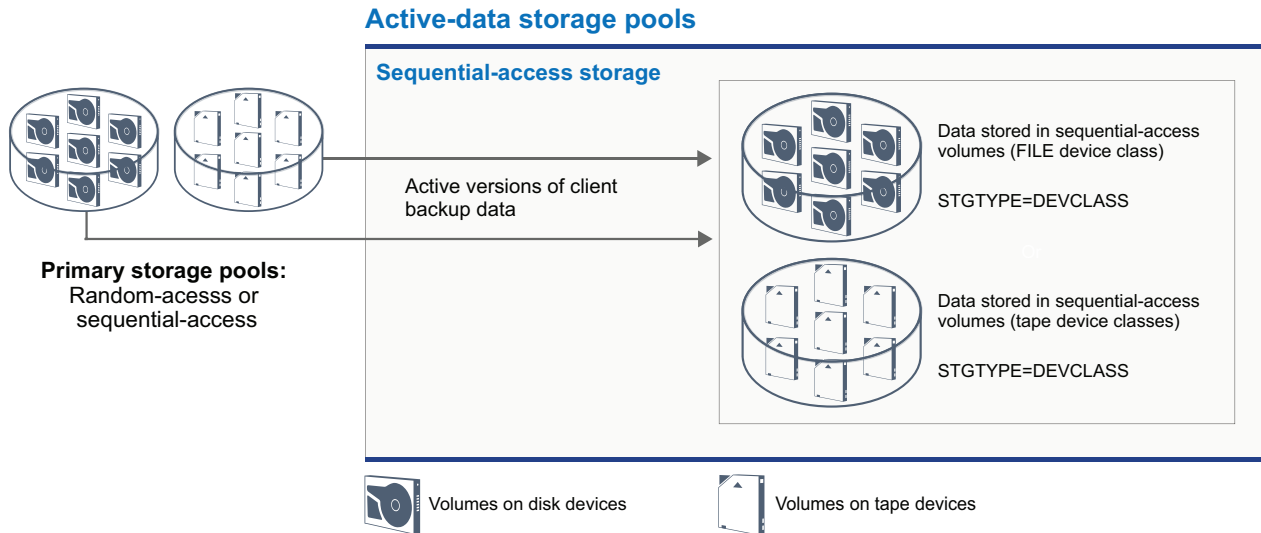


Figure 9. Active-data storage pools

Active-data pools can use any type of sequential-access storage. However, the benefits of an active-data pool depend on the device type that is associated with the pool. For example, active-data pools that are associated with a FILE device class are ideal for fast client restore operations because of the following reasons:

- FILE volumes do not have to be physically mounted
- Client sessions that are restoring from FILE volumes in an active-data pool can access the volumes concurrently, which improves restore performance

Data transport to storage across networks

The Tivoli Storage Manager environment provides ways to securely move data to storage across various types of networks and configurations.

Network configurations for storage devices

Tivoli Storage Manager provides methods for configuring clients and servers on a local area network (LAN), on a storage area network (SAN), LAN-free data movement, and as network-attached storage.

Data backup operations over a LAN

Figure 10 on page 23 shows the data path for Tivoli Storage Manager backup operations over a LAN.




Figure 10. Tivoli Storage Manager backup operations over a LAN

In a LAN configuration, one or more tape libraries are associated with a single Tivoli Storage Manager server. In this type of configuration, client data, electronic mail, terminal connection, application program, and device control information must all be handled by the same network. Device control information and client backup and restore data flow across the LAN.

Data backup operations over a SAN

Figure 11 shows the data path for Tivoli Storage Manager backup operations over a SAN.




Figure 11. Tivoli Storage Manager backup operations over a SAN

A SAN is a dedicated storage network that can improve system performance. On a SAN, you can consolidate storage and relieve the distance, scalability, and bandwidth limitations of LANs and wide area networks (WANs). By using Tivoli Storage Manager in a SAN, you can take advantage of the following functions:

- Share storage devices among multiple Tivoli Storage Manager servers. Devices that use the GENERICTAPE device type are not included.
- Move data from a client system directly to storage devices without using the LAN. LAN-free data movement requires the installation of a storage agent on the client system. The storage agent is available with the Tivoli Storage Manager for Storage Area Networks product.

Through the storage agent, the client can directly back up and restore data to a tape library or shared file system such as GPFS™. The Tivoli Storage Manager server maintains the server 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 bandwidth on the LAN that would otherwise be used for client data movement.

- Share tape drives and libraries that are supported by the Tivoli Storage Manager server.
- Consolidate multiple clients under a single client node name in a General Parallel File System (GPFS) cluster.

Network-attached storage

Network-attached storage (NAS) file servers are dedicated storage servers whose operating systems are optimized for file-serving functions. NAS file servers typically interact with Tivoli Storage Manager through industry-standard network protocols, such as network data management protocol (NDMP) or as primary storage for random-access or sequential access storage pools. Tivoli Storage Manager provides the following basic types of configurations that use NDMP for backing up and managing NAS file servers:

- Tivoli Storage Manager backs up a NAS file server to a library device that is directly attached to the NAS file server. The NAS file server, which can be remote from the Tivoli Storage Manager server, transfers backup data directly to a drive in a SCSI-attached tape library. Data is stored in NDMP-formatted storage pools, which can be backed up to storage media that can be moved offsite for protection in case of an onsite disaster.
- Tivoli Storage Manager backs up a NAS file server over the LAN to a storage-pool hierarchy. In this type of configuration, you can store NAS data directly to disk, either random access or sequential access, and then migrate the data to tape. You can also use this type of configuration for system replication. Data can also be backed up to storage media that can be moved offsite. The advantage of this type of configuration is that you have all of the data management features associated with a storage pool hierarchy.
- The Tivoli Storage Manager client reads the data from the NAS system by using NFS or CIFS protocols and sends the data to the server to be stored.

Storage management

You manage the devices and media that are used to store client data through the Tivoli Storage Manager server. The server integrates storage management with the policies that you define for managing client data in the following areas:

Types of devices for server storage

With Tivoli Storage Manager, you can use directly attached devices and network-attached devices for server storage. Tivoli Storage Manager represents physical storage devices and media with administrator-defined storage objects.

Data migration through the storage hierarchy

For primary storage pools other than directory-container storage pools, you can organize the storage pools into one or more hierarchical structures. This storage hierarchy provides flexibility in a number of ways. For example, you can set a policy to back up data to disks for faster backup operations. The Tivoli Storage Manager server can then automatically migrate data from disk to tape.

Removal of expired data

The policy that you define controls when client data automatically expires from the Tivoli Storage Manager server. To remove data that is eligible for expiration, a server expiration process marks data as expired and deletes metadata for the expired data from the database. The space that is occupied by the expired data is then available for new data. You can control the frequency of the expiration process by using a server option.

Media reuse by reclamation

As server policies automatically expire data, the media where the data is stored accumulates unused space. For storage media other than directory-container storage pools or random disk storage pools, the Tivoli Storage Manager server implements *reclamation*, a process that frees media for reuse without traditional tape rotation. Reclamation automatically defragments media by consolidating unexpired data onto other media when the free space on media reaches a defined level. The reclaimed media can then be used again by the server. Reclamation allows media to be automatically circulated through the storage management process and minimize the number of media that are required.

Consolidating backed up client data

By grouping the client data that is backed up, you can minimize the number of media mounts required for client recovery. The Tivoli Storage Manager server provides the following methods for grouping client files on storage media other than directory-container storage pools:

Collocating client data

The Tivoli Storage Manager server can *collocate* client data, in other words store client data on a few volumes instead of spreading the data across many volumes. Collocation by client minimizes the number of volumes that are required to back up and restore client data. Data collocation might increase the number of volume mounts because each client might have a dedicated volume instead of data storage from several clients in the same volume.

You can set the server to collocate client data when the data is initially placed in server storage. In a storage hierarchy, you can collocate the data when the server migrates the data from the initial storage pool to the next storage pool in the storage hierarchy. You can collocate by client, by file space per client, or by a group of clients. Your selection depends on the size of the file spaces that are stored and restore requirements.

Associating active-data pools with various devices

Active-data pools are useful for fast restoration of client data. Benefits include a reduction in the number of onsite or offsite storage volumes, or reducing bandwidth when you copy or restore files that are vaulted electronically in a remote location. Active-data pools that use removable media, such as tape, offer similar benefits. Although tape devices must be mounted, the server does not have to position past inactive files. However, the primary benefit of using removable media in active-data pools is that the number of volumes that are used for onsite and offsite storage is reduced. If you store data to a remote location, you can minimize the amount of data that must be transferred by copying and restoring only active data.

Creating a backup set

A backup set contains all of the active backed-up files that exist for that client in server storage. The backup set is portable and is retained for the time that you specify. A backup set is in addition to the backups that are already stored and requires extra media.

Moving data for a client node

You can consolidate data for a client node by moving the data within server storage. You can move a backup set to different media, where the backup set is retained until the time that you specify. Consolidating data can help to improve efficiency during client restore or retrieve operations.

Chapter 3. Data protection strategies with Tivoli Storage Manager

Tivoli Storage Manager provides ways for you to implement various data protection strategies.

You can configure Tivoli Storage Manager to send data to storage devices that are on the local site or on a remote site. To maximize data protection, you can configure replication to a remote server.

Strategies to minimize the use of storage space for backups

To minimize the amount of storage space that is required, Tivoli Storage Manager backs up data by using the data deduplication and progressive incremental backup techniques.

Data deduplication

When the Tivoli Storage Manager server receives data from a client, the server identifies duplicate data extents and stores unique instances of the data extents in a directory-container storage pool. The data deduplication technique improves storage utilization and eliminates the need for a dedicated data deduplication appliance.




Figure 12. Data deduplication process

If the same byte pattern occurs many times, data deduplication greatly reduces the amount of data that must be stored or transferred. In addition to whole files, Tivoli Storage Manager can also deduplicate parts of files that are common with parts of other files.

Tivoli Storage Manager provides the following types of data deduplication:

Server-side data deduplication

The server identifies duplicate data extents and moves the data to a directory-container storage pool. The server-side process uses *inline data*

deduplication, where data is deduplicated at the same time that the data is written to a directory-container storage pool. Deduplicated data can also be stored in other types of storage pools. Inline data deduplication on the server provides the following benefits:

- Eliminates the need for reclamation
- Reduces the space that is occupied by the stored data

Client-side data deduplication

With this method, processing is distributed between the server and the client during a backup process. The client and the server identify and remove duplicate data to save storage space on the server. In client-side data deduplication, only compressed, deduplicated data is sent to the server. The server stores the data in the compressed format that is provided by the client. Client-side data deduplication provides the following benefits:

- Reduces the amount of data that is sent over the local area network (LAN)
- Eliminates extra processing power and time that is required to remove duplicate data on the server
- Improves database performance because the client-side data deduplication is also inline

You can combine both client-side and server-side data deduplication in the same production environment. The ability to deduplicate data on either the client or the server provides flexibility in terms of resource utilization, policy management, and data protection.

Compression

Use inline compression to reduce the amount of space that is stored in container storage pools. Data is compressed as it is written to the container storage pool.

Restriction: The Tivoli Storage Manager server cannot compress encrypted data.

Progressive incremental backup

In a progressive incremental backup process, the server monitors client activity and backs up any files that change since the initial full backup. Entire files are backed up, so that the server does not need to reference base versions of the files. This backup technique eliminates the need for multiple full backups of client data thus saving network resources and storage space.

Strategies for disaster protection

Tivoli Storage Manager provides strategies to protect data if a disaster occurs. These strategies include node replication to a remote site, storage pool protection, database protection, and device replication to a standby server.

Replication to a remote site

Replication is the process of incrementally copying data from one server to a remote server. The server from which client data is replicated is called a *source replication server*. The server to which client data is replicated is called a *target replication server*. A replication server can function as a source server, a target server, or both. You use replication processing to maintain the same level of files on the source and

the target servers. Figure 13 shows the node replication process to a remote site.




Figure 13. Node replication process

When client data is replicated, data that is not on the target server is copied. When replicated data exceeds the retention limit, the target server automatically removes the data. To maximize data protection, you synchronize the local server and the remote server, for example Site B replicates data from Site A and Site A replicates data from Site B. As part of replication processing, client data that was deleted from the source server is also deleted from the target server.

Tivoli Storage Manager provides the following replication functions:

- You can define policies for the target server for the following conditions:
 - Identical policies to the source server
 - Different policies to cater for different business requirements

If a disaster occurs and the source server is unavailable, clients can recover data from the target server. If the source server cannot be recovered, you can convert clients to store data on the target server. When an outage occurs, the source server can automatically fail over to a target server for data recovery.

- You can use replication processing to recover damaged files from FILE or DISK storage pools. You must replicate the client data to the target server before the file damage occurs. Subsequent replication processes detect damaged files on the source server and replace the files with undamaged files from the target server.

Role of replication in disaster protection

If a disaster occurs, you can recover replicated data from the remote site and maintain the same level of files on the source and target servers. You use replication to achieve the following objectives:

- Control network throughput by scheduling node replication at specific times
- Recover data from large-scale site loss

- Recover damaged files on the source server

To ensure a high level of server availability, you can use replication with a group of servers and clients, in other words *clustering*. In a clustered environment, a client is less likely to fail over to another server. If you replicate data from several source servers to one target server, you create a high dependency on the target server. A clustered environment reduces the dependency on the target server.

Storage pool protection

You use different techniques to protect against the permanent loss of data that is stored in directory-container storage pools and in FILE and DISK storage pools.

Directory-container storage pools

You use storage pool protection to protect specific directory-container storage pools if you do not need to replicate all the data that is contained in a client node. By protecting a directory-container storage pool, you do not use resources that replicate existing data and metadata, which improves server performance.

For best data protection, protect the directory-container storage pool before you replicate the client node. When node replication is started, the data extents that are already replicated through storage pool protection are skipped, which reduces the replication processing time. You can repair specific directory-container storage pools.

Storage pools that are associated with FILE and DISK device classes

You use node replication to protect storage pools that are associated with FILE and DISK device classes. When you restore a storage pool, the server determines which files are in that storage pool. Using file copies from a copy storage pool or an active-data storage pool, the server restores the files to the same or a different storage pool. As part of a restore operation, inactive file versions are deleted from the server database. Inactive file versions are deleted if the server determines that an inactive file needs to be replaced but cannot find the file in the active-data pool.

Database protection

You can use the replication technique to protect against the permanent loss of data that is stored in the server database. The database backup can be full, incremental, or snapshot. To restore the database, you must have the backup volumes for the database. You can restore the database from backup volumes by either a point-in-time restore or a most current restore operation.

Point-in-time restore

The point-in-time restore operation includes the following actions:

- Removes and re-creates the active log directory and archive log directory that is specified in `dsmserv.opt` file.
- Restores the database image from backup volumes to the database directories recorded in a database backup or to new directories.
- Restores archive logs from backup volumes to the overflow directory.
- Applies logs from the overflow directory up to a specified point in time.

Most current restore

The most current restore operation includes the following actions:

- Does not remove and re-create the active log directory or archive log directory.
- Restores a database image from the backup volumes to the database directories recorded in a database backup or to new directories.
- Restores archive logs from backup volumes to the overflow directory.
- Applies logs from overflow directory and archive logs from archive log directory.

Use point-in-time restore operations for situations such as disaster recovery or to remove the effects of errors that can cause inconsistencies in the database. Restore operations for the database that use snapshot backups are a form of point-in-time restore operation. If you want to recover the database to the time when the database was lost, recover the database to the most current state.

Alternative methods for disaster protection

In addition to replication to a remote site, you can also use the following methods to protect data and implement disaster recovery with Tivoli Storage Manager:

Sending backup tapes to a remote site

Data is backed up to tape at scheduled times by the source server. The tapes are sent to a remote site. If a disaster occurs, the tapes are returned to the site of the source server and the data is restored on the source clients.

Multisite appliance replication to a standby server

In the multisite appliance configuration, the source appliance is replicated to a remote server in a SAN architecture. In this configuration, if the client hardware at the original site is damaged or destroyed, the source device can be replicated from the standby server at the remote site. This configuration provides disk-based backup and restore operations.

Strategies for disaster recovery with Tivoli Storage Manager

Tivoli Storage Manager provides several ways to recover the server if the database or storage pools fail.

Automatic failover for disaster recovery

Automatic failover is an operation that switches to a standby system if a software, hardware, or network interruption occurs. Automatic failover is used with node replication to recover data after a system failure. Figure 14 on page 32 shows the Tivoli Storage Manager automatic failover process.

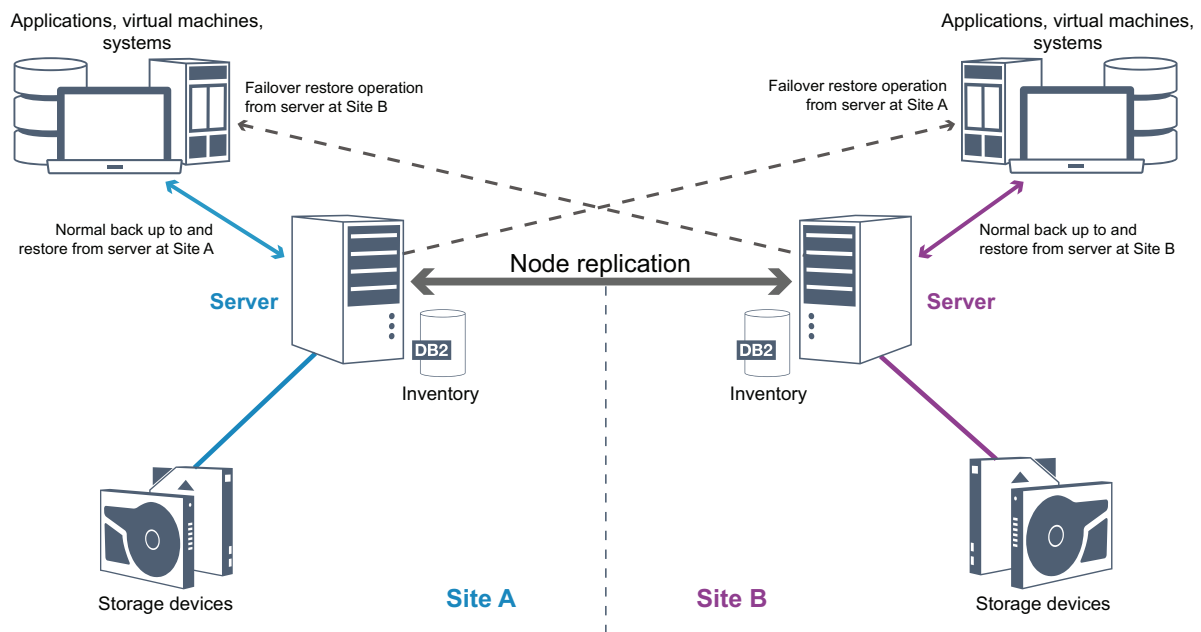


Figure 14. Automatic failover process

Automatic failover for data recovery occurs if the source replication server is unavailable because of a disaster or a system outage. During normal operations, when the client accesses a source replication server, the client receives connection information for the target replication server. The client node stores the failover connection information in the client options file.

During client restore operations, the server automatically changes clients from the source replication server to the target replication server and back again. Only one server per node can be used for failover protection at any time. When a new client operation is started, the client attempts to connect to the source replication server. The client resumes operations on the source server if the source replication server is available.

To use automatic failover for replicated client nodes, the source replication server, the target replication server, and the client must be at the Tivoli Storage Manager V7.1 level or later. If any of the servers are at an earlier level, automatic failover is disabled and you must rely on a manual failover process.

Recovery of Tivoli Storage Manager components

The server database, recovery log, and storage pools are critical to the operation of Tivoli Storage Manager and must be protected. If the database is unusable, the entire server is unavailable and recovering data that is managed by the server might be difficult or impossible.

Even without the database, fragments of data or complete files might be read from storage pool volumes that are not encrypted and security can be compromised. Therefore, you must always back up the database. Also, always encrypt sensitive data by using the client or the storage device, unless the storage media is physically secured.

Tivoli Storage Manager provides several data protection methods, which include backing up storage pools and the database. For example, you can define schedules so that the following operations occur:

- After the initial full backup of your storage pools, incremental storage pool backups are run every night.
- Incremental database backups are run every night.
- Full database backups are run once a week.

For tape-based environments, you can use disaster recovery manager (DRM) to assist you in many of the tasks that are associated with protecting and recovering data. DRM is available with Tivoli Storage Manager Extended Edition.

Preventive actions for recovery

Recovery is based on the following preventive actions:

- Mirroring, by which the server maintains a copy of the active log
- Backing up the database
- Backing up the storage pools
- Auditing storage pools for damaged files and recovery of damaged files when necessary
- Backing up the device configuration and volume history files
- Validating the data in storage pools by using cyclic redundancy checking
- Storing the `cert.kdb` file in a safe place to ensure that the Secure Sockets Layer (SSL) is secure

If you are using tape for storage, you can also create a disaster recovery plan to guide you through the recovery process by using DRM. You can use the disaster recovery plan for audit purposes to certify the recoverability of the server. The disaster recovery methods of DRM are based on taking the following actions:


- Creating a disaster recovery plan file for the server
- Backing up server data to tape
- Sending the server backup data to a remote site or to another server
- Storing client system information
- Defining and tracking the storage media that is used for storing and recovering client data

Part 2. Tivoli Storage Manager solutions for data protection

To help you to deploy a data protection environment, review information about Tivoli Storage Manager configurations, and select the best solution for your business needs.

Chapter 4. Disk-based implementation of a data protection solution for a single site

This disk-based implementation of a data protection solution with Tivoli Storage Manager uses inline data deduplication and provides protection for data on a single site.



This data protection solution provides the following benefits:

- Server system and storage hardware at a single site
- Cost-effective use of storage through the data deduplication feature
- Space-efficient solution with minimal hardware setup
- Minimal implementation that requires installation and configuration for only one server and supporting storage hardware

In this solution, the client sends data to the Tivoli Storage Manager server, where the data is deduplicated and stored in a directory-container storage pool that is implemented in disk storage. Data from the inventory is also backed up to disk storage. This solution is suitable for entry-level environments for which a second copy of data is not required.


Related reference:

Chapter 8, "Comparison of data protection solutions," on page 45

Chapter 9, "Roadmap for implementing a data protection solution," on page 47

Chapter 5. Disk-based implementation of a data protection solution for multiple sites

This disk-based implementation of a data protection solution with Tivoli Storage Manager uses inline data deduplication and replication at two sites.



This data protection solution provides the following benefits:

- Replication can be configured at both sites so that each server protects data for the other site
- Offsite data storage for each location is simplified
- Bandwidth is used efficiently because only deduplicated data is replicated between the sites
- Clients can automatically fail over to a target replication server if the source replication server is unavailable

In this solution, clients send data to the source server, where the data is deduplicated and stored in a directory-container storage pool that is implemented in disk storage. The data is replicated to the storage pool on the target server for each site. This solution is suitable for environments that require disaster protection. If mutual replication is configured, clients at both sites can use failover recovery for continued backups and data recovery from the available server on the other site.

Related reference:

Chapter 8, "Comparison of data protection solutions," on page 45

Chapter 9, "Roadmap for implementing a data protection solution," on page 47

Chapter 6. Tape-based implementation of a data protection solution

This implementation of a data protection solution with Tivoli Storage Manager uses one or more tape storage devices to back up data. Tape backup provides low-cost scalability that is optimized for long-term retention.

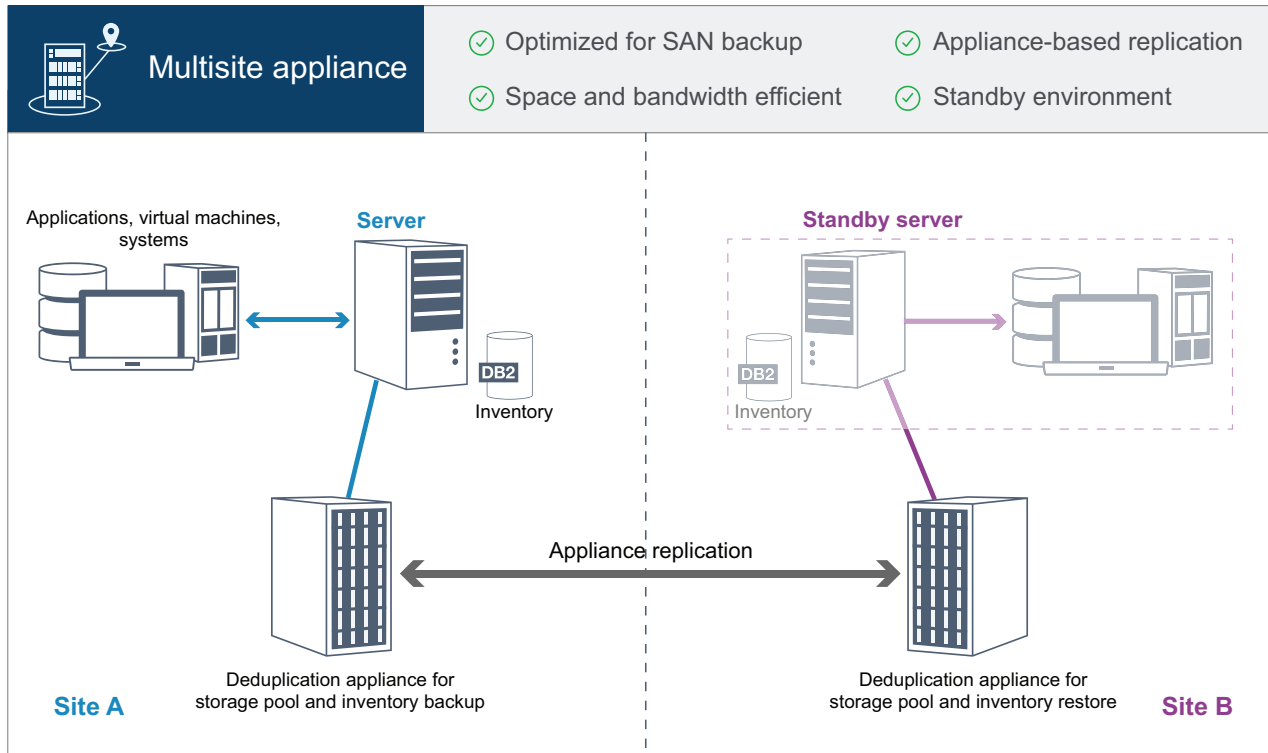
Related reference:

Chapter 8, “Comparison of data protection solutions,” on page 45

Chapter 9, “Roadmap for implementing a data protection solution,” on page 47

Chapter 7. Appliance-based implementation of a data protection solution for multiple sites

This implementation of a multi-site Tivoli Storage Manager data protection solution uses appliance-based data deduplication and replication. A standby server is configured at a second site to recover data if the primary server is unavailable.



This data protection solution provides the following benefits:

- Performance is optimized for backups on high-speed storage area networks (SAN) and for use with Tivoli Storage Manager for SAN, when clients back up directly to SAN-attached virtual tape devices.
- Fast, appliance-based replication frees the server from having to track replication metadata in the server database.
- Bandwidth and storage space are used efficiently because only deduplicated data is replicated between the sites.
- A standby environment provides for disaster recovery, but does not require the amount of resources that are needed for a fully active site.

In this data protection configuration, the server uses hardware appliances to deduplicate and replicate data. The appliance at Site A deduplicates data and then replicates the data to the appliance at Site B for disaster protection. If a failure at Site A occurs, you make the standby server active by restoring the most recent database backup, and by activating the replicated copy of data.





Related reference:





Chapter 8, “Comparison of data protection solutions,” on page 45

Chapter 9, "Roadmap for implementing a data protection solution," on page 47

Chapter 8. Comparison of data protection solutions

Compare the key features for each Tivoli Storage Manager solution to determine which configuration best meets your data protection requirements. Then, review the available documentation to implement the solution.

	Single-site disk	Multisite disk	Multisite appliance	Tape
				
Highlights				
Cost	\$	\$\$\$	\$\$\$\$	\$\$
Protection level	One data copy	Two or more data copies	Two or more data copies	Two or more data copies
Disaster recovery	None	Active server	Standby server	Offsite copies
Key benefits				
Leading-edge data reduction	✓	✓	✓	✓
Fast and efficient disk-based backup and restore operations	✓	✓	✓	
Simplified offsite management		✓		
Data deduplication feature at no extra cost	✓	✓		
Replication processing included at no extra charge		✓		
Data deduplication at both the source and target server		✓		
Low-cost scalability and optimized for long-term retention				✓
Efficiency and cost				
Optimized for high-speed storage area network (SAN) backup operations			✓	✓
Optimized for high-speed local area network (LAN)	✓	✓	✓	
Global data deduplication across all data types and sources	✓	✓	✓	
Bandwidth-efficient replication		✓	✓	
Lower energy costs				✓
Option for a second copy without extra disk hardware				✓
Availability				
Offsite copy capability		✓	✓	✓
Appliance-based replication			✓	

	Single-site disk	Multisite disk	Multisite appliance	Tape
				
Client recovery from high-availability server		✓		
Replication target in the cloud		✓		
Independent management of retention policies for replication data; ability to keep more or less data at recovery site		✓		
Application-level replication; ability to choose which systems and applications are replicated		✓		
Scalability				
Global data deduplication across servers			✓	
SAN-optimized backup directly to tape for large data types				✓
Single-instance petabyte scalability				✓

What to do next

Review available documentation for the solutions in Chapter 9, “Roadmap for implementing a data protection solution,” on page 47.

Related reference:

Chapter 4, “Disk-based implementation of a data protection solution for a single site,” on page 37

Chapter 5, “Disk-based implementation of a data protection solution for multiple sites,” on page 39

Chapter 7, “Appliance-based implementation of a data protection solution for multiple sites,” on page 43

Chapter 6, “Tape-based implementation of a data protection solution,” on page 41

Chapter 9. Roadmap for implementing a data protection solution

Plan and implement the most suitable data protection solution for your business environment with Tivoli Storage Manager.

Single-site disk solution

For steps that describe how to plan for, implement, monitor, and operate a single-site disk solution, see [Single-site disk solution](#).

Multisite disk solution

For steps that describe how to plan for, implement, monitor, and operate a multisite disk solution, see [Multisite disk solution](#).

Tape solution

For a high-level overview of the tasks that are required to implement a tape solution, review the following steps:

1. Begin planning for the solution by reviewing V7.1.1 information at the following links:
 - [Planning for server storage \(V7.1.1\)](#)
 - [Capacity planning \(V7.1.1\)](#)
2. Install the server and optionally, the Operations Center. Review information at the following links:
 - [Installing the server](#)
 - [Installing and upgrading the Operations Center](#)
3. Configure the server for tape storage. Review V7.1.1 information at the following links:
 - [AIX](#) [Linux](#) [Configuring and managing storage devices \(V7.1.1\)](#)
 - [Windows](#) [Configuring and managing storage devices \(V7.1.1\)](#)

For guidance about performance best practices, see [Configuration best practices](#).

4. Configure policies to protect your data. Review the information in [Customizing policies](#).
5. Set up client schedules. Review the information in [Scheduling operations for client nodes \(V7.1.1\)](#).
6. Install and configure clients. To determine the type of client software that you require, review the information in [Adding clients for details](#).
7. Configure monitoring for your system. Review the information in [Monitoring storage solutions](#).

Multisite appliance solution

For an overview of the tasks that are required to implement a multisite appliance solution, review the following steps:

1. Begin planning for the solution by reviewing V7.1.1 information at the following links:

- Planning for server storage (V7.1.1)
 - Capacity planning (V7.1.1)
2. Install the server and optionally, the Operations Center. Review information at the following links:
 - Installing the server
 - Installing and upgrading the Operations Center
 3. Configure the server for storage in a virtual tape library. Review V7.1.1 information at the following links:
 - AIX Linux Configuring and managing storage devices (V7.1.1)
 - Windows Configuring and managing storage devices (V7.1.1)
 - Managing virtual tape libraries (V7.1.1)

For guidance about improving system performance, see Configuration best practices.

4. Configure policies to protect your data. Review the information in Customizing policies.
5. Set up client schedules. Review the information in Scheduling backup and archive operations.
6. Install and configure clients. To determine the type of client software that you need, review the information in Adding clients for details.
7. Configure monitoring for your system. Review the information in Monitoring storage solutions.

Related reference:

Chapter 8, “Comparison of data protection solutions,” on page 45

Chapter 4, “Disk-based implementation of a data protection solution for a single site,” on page 37

Chapter 5, “Disk-based implementation of a data protection solution for multiple sites,” on page 39

Chapter 7, “Appliance-based implementation of a data protection solution for multiple sites,” on page 43

Chapter 6, “Tape-based implementation of a data protection solution,” on page 41

Part 3. Appendixes

Appendix. Accessibility features for the Tivoli Storage Manager product family

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

Accessibility features

The IBM Tivoli Storage Manager family of products includes the following accessibility features:

- Keyboard-only operation using standard operating-system conventions
- Interfaces that support assistive technology such as screen readers

The command-line interfaces of all products in the product family are accessible.

Tivoli Storage Manager Operations Center provides the following additional accessibility features when you use it with a Mozilla Firefox browser on a Microsoft Windows system:

- Screen magnifiers and content zooming
- High contrast mode

The Operations Center and the Tivoli Storage Manager server can be installed in console mode, which is accessible.

The Operations Center help system is enabled for accessibility. For more information, click the question mark icon on the help system menu bar.

Vendor software

The Tivoli Storage Manager product family includes certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for the accessibility information about its products.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (<http://www.ibm.com/able>) for information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the U.S.A.

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

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

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

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

Portions of this code are derived from IBM® Corp. Sample Programs.

© Copyright IBM® Corp. _enter the year or years_. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights Except as expressly granted in this permission, no other permissions,

licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

A glossary is available with terms and definitions for the IBM Tivoli Storage Manager family of products.

See Tivoli Storage Manager glossary.

To view glossaries for other IBM products, see IBM Terminology.

Index

A

- About this publication v
- accessibility features 51
- active-data pools 22
- active-data storage pools 17
- API
 - See application programming interface
- application clients 4
- application programming interface 10
- archive service 4

B

- backup service 4

C

- client data
 - consolidation of 22
 - create a backup set for 22
 - management of 22
 - migration of 22
 - moving to storage 22
- clients
 - applications 4
 - client nodes 3
 - client software 3
 - concepts 3
 - system clients 4
 - types of 4
 - virtual machines 4
- cloud-container storage pools 17
- collocation 22
- command-line interface 10
- concepts
 - clients 3
 - database 3
 - inventory 3
 - overview 3
 - recovery log 3
 - server 3
 - storage 3
- container storage pools 27
- container-copy storage pools 17
- copy storage pools 17

D

- data deduplication
 - client-side 27
 - inline 27
 - server-side 27
- data mover 13
- data protection
 - strategies 27
- data protection services 4
- device class 13
- device replication 28, 31
- directory-container storage pools 17

- disability 51
- disaster recovery
 - automatic failover 31
 - DRM 31
 - manager 31
 - methods 28
 - preventive measures 31
- drive 13

F

- failover, automatic 31

G

- GUI, for clients 10

I

- IBM Knowledge Center v
- inline data deduplication 27
- interfaces
 - API 10
 - backup-archive client 10
 - client GUI 10
 - command-line 10
 - operations center 10
 - SQL statements 10
- inventory 7

K

- keyboard 51
- Knowledge Center v

L

- layer
 - logical 13
 - physical 13
- library 13
- log
 - active log 7
 - archive failover log 7
 - archive log 7
 - log mirror 7
 - recovery log 7

M

- media
 - reclamation of 22
- media, removable 13
- migrate service 4

N

- network, types of
 - LAN 22
 - LAN-free 22
 - NAS 22
 - Network attached storage 22
 - SAN 22
- node replication 28, 31

O

- operating systems 4
- operations center
 - access to 10
 - functions 10

P

- path 13
- policy
 - data management by 7
 - policy domain 7
 - policy set 7
 - standard 7
- primary storage pools 17
- progressive incremental backup 27
- publications v

R

- recall service 4
- recovery
 - data 31
 - system components 31
- recovery log 7
- replication
 - node 28
 - role in disaster recovery 28
 - source server 28
 - target server 28
- restore service 4
- retrieve service 4

S

- SAN architecture 28, 31
- security management
 - closed registration 7, 22
 - open registration 7, 22
 - passwords 7, 22
 - SSL 7, 22
 - TLS 7, 22
- server 13
 - concepts 3
 - data stores 7
 - inventory 7
 - recovery log 7
- services
 - archive and retrieve 4
 - backup and restore 4
 - migrate and recall 4
- solutions
 - data protection solutions
 - appliance-based 43

- solutions (*continued*)
 - multisite solution
 - appliance-based 43
- SQL statements, to access server database 10
- storage
 - concepts 3
 - device support for 22
 - devices 3, 13
 - hierarchy 3, 22
 - management of 22
 - networks 22
 - objects 13
 - pools 3, 13, 17
 - representations 13
 - types 13
 - volumes 17
- storage pools
 - archive-data 17
 - cloud 17
 - container 17, 27
 - container-copy 17
 - copy 17
 - primary 17
 - representation 17
 - types of 17
- system clients 4

T

- tape devices
 - physical 13
 - virtual 13
- tape transport 28, 31
- Tivoli Storage Manager solutions
 - data protection solutions 41
 - comparison 45
 - multisite disk 39
 - single-site disk 37
 - multisite solution
 - disk-based 39
 - roadmap 47
 - single-site solution
 - disk-based 37

V

- virtual machines 4
- volume 13
- volumes 17

W

- web interface, for backup-archive client 10



Product Number: 5608-E01
5608-E02
5608-E03

Printed in USA