DDOS 7.12

Administration Guide

7.12

July 2023
Rev. 02

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Introducing DDOS

This chapter presents the following topics:

Topics:

Revision history

The revision history lists the major changes to this document.

Table 1. Document revision history

Revision	Date	Description
02 (7.12)	July 2023	Removed references to the DS600 disk shelf.
01 (7.12)	July 2023	This revision includes information about: • Clarified groupname requirements when using LDAP authentication for Active Directory. • Updated limitations on physical interface bonding. • Updated default behavior for MTree analytics. • Updated upgrade behavior with MTree analytics. • New AWS regions for Cloud Tier. • Retention Lock Compliance Support on cloud-based DDVE instances.

System overview

Dell PowerProtect DD Series Appliances and older Data Domain systems are disk-based appliances that run PowerProtect DDOS to provide inline deduplication for data protection and disaster recovery (DR) in the enterprise environment.

NOTE: In this guide, "DD system," "the protection system," or simply "the system" refers to PowerProtect DD Series Appliances running DDOS 7.0 or later as well as earlier Data Domain systems.

DD system appliances vary in storage capacity and data throughput. Systems are typically configured with expansion enclosures that add storage space.

DDOS provides the following interfaces:

- DD System Manager—Enables you to configure, manage, and monitor your system using a browser-based graphical user interface (GUI). DD System Manager provides real-time graphs and tables that enable you to monitor the status of system hardware components and configured features. DD System Manager provides a single, consolidated management interface that enables you to manage a single system from any location.

  NOTE: If you have a larger environment, PowerProtect DD Management Center (DDMC) enables you to manage multiple systems from a single browser window. Contact your Dell representative for more information.
- Command-line interface (CLI)—Enables you to perform all system operations, including operations that cannot be managed by DD System Manager. Using the CLI commands you can configure system settings and display system hardware status, feature configuration, and operations. Refer to the *PowerProtect DD Series Appliances Operating System Command Reference Guide* for a complete description of commands.

DDOS features

DDOS features include:

- Data integrity—The DDOS Data Invulnerability Architecture protects against data loss from hardware and software failures.
- Data Deduplication—The file system deduplicates data by identifying redundant data during each backup and storing unique data just once.
- Restore operations—File restore operations create little or no contention with backup or other restore operations.
- DD Replicator—DD Replicator sets up and manages the replication of backup data between two protection systems.
- Multipath and load balancing—In a Fibre Channel multipath configuration, multiple paths are established between a protection system and a backup server or backup destination array. When multiple paths are present, the system automatically balances the backup load between the available paths.
- High availability—The High Availability (HA) feature lets you configure two protection systems as an Active-Standby pair, providing redundancy in the event of a system failure. HA keeps the active and standby systems in sync, so that if the active node were to fail due to hardware or software issues, the standby node can take over services and continue where the failing node left off.
- Random I/O handling—The random I/O optimizations included in DDOS provide improved performance for applications and use cases that generate larger amounts of random read and write operations than sequential read and write operations.
- System Administrator access—System administrators can access the system for configuration and management using a command line interface (CLI) or a graphical user interface (GUI).
- Licensed features—Feature licenses allow you to purchase only those features you intend to use. Some examples of features that require licenses are DD Boost, and capacity on demand (storage capacity increases).
- Storage environment integration—DDOS systems integrate easily into existing data centers.

Licensed features

Feature licenses allow you to purchase only those features you intend to use. Some examples of features that require licenses are DD Boost, and capacity on demand (storage capacity increases).

Consult with your sales representative for information on purchasing licensed features.

Table 2. Features requiring licenses

Feature Name	License Name in Software	Description
DD ArchiveStore	ARCHIVESTORE	Licenses systems for archive use, such as file and email archiving, file tiering, and content and database archiving.
DD Boost	DDBOOST	Enables the use of a system with qualified backup software. The online compatibility guide available from E-Lab Navigator provides the list of qualified applications. The managed file replication (MFR) feature of DD Boost also requires the DD Replicator license.
DD Capacity on Demand	CONTROLLER-COD	Enables an on-demand capacity increase for a DD system that is not at its maximum supported capacity.
Cloud Tier	CLOUDTIER-CAPACITY	Enables a system to move data from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention. This type of licenses does not apply to DD6400 systems, which use a Protection Pool license.
DD Encryption	ENCRYPTION	Allows data on system drives or external storage to be encrypted while being saved and locked when moving the system to another location.
DD Expansion Storage	EXPANDED-STORAGE	Allows system storage to be expanded beyond the level provided in the base system.
DD I/OS (for IBM i operating environments)	I/OS	An I/OS license is required when DD VTL is used to backup systems in the IBM i operating environment. Apply this license before adding virtual tape drives to libraries.
DD Replicator	REPLICATION	Adds DD Replicator for replication of data from one protection system to another. A license is required on each system.

Table 2. Features requiring licenses (continued)

Feature Name	License Name in Software	Description
DD Retention Lock Compliance Edition	RETENTION-LOCK-COMPLIANCE	When properly installed, configured, enabled, and administered, enables you to meet strict data permanence requirements of regulatory standards, such as those of SEC 17a-4(f).
DD Retention Lock Governance Edition	RETENTION-LOCK-GOVERNANCE	Protects selected files from modification and deletion before a specified retention period expires.
DD Shelf Capacity-Active Tier	CAPACITY-ACTIVE	Enables a system to expand the active tier storage capacity to an additional enclosure or a disk pack within an enclosure. This type of licenses does not apply to DD6400 systems, which use a Protection Pool license.
DD Storage Migration	STORAGE-MIGRATION-FOR-DATADOMAIN-SYSTEMS	Enables migration of data from one enclosure to another to support replacement of older, lower-capacity enclosures.
DD Virtual Tape Library (DD VTL)	VTL	Enables the use of a protection system as a virtual tape library over a Fibre Channel network. This license also includes the NDMP Tape Server feature and the I/OS license for IBM i systems, which previously required separate licenses.
High Availability	HA-ACTIVE-PASSIVE	Enables the High Availability feature in an Active-Standby configuration. You only need to purchase one HA license; the license runs on the active node and is mirrored to the standby node.
Protection Pool	PROTECTION-POOL	Enables Active Tier, Cloud Tier, and Cache Tier storage on DD6400 systems.
SSD Cache	SSD-CAPACITY	Enables the SSD cache feature on DD6300, DD6800, DD9300, and DD9800 systems. This license is not required to use the SSD cache feature on DD6900, DD9400, and DD9900 systems.

2

Getting Started

This chapter presents the following topics:

Topics:

Security updates

Attention: There may be security updates available for this product. Please always consult https://www.dell.com/support/security, prior to commissioning, to determine if any applicable security updates are available for download and application. We further encourage you to sign up for Security Alerts at the same URL to be proactively notified of new Security Alerts that Dell issues.

There may be security updates available for this product. Please always consult https://www.dell.com/support/security, prior to commissioning, to determine if any applicable security updates are available for download and application. We further encourage you to sign up for Security Alerts at the same URL to be proactively notified of new Security Alerts that Dell issues.

Logging in and out of DD System Manager

Prerequisites

DD System Manager is reachable from any TCP port. You can allow or disallow specific port numbers according to your security requirements.

About this task

When you connect to DD System Manager from a web browser, all HTTP connections are automatically redirected to HTTPS.

Use one of the following methods to log in to DD System Manager:

- Log in with a user name and password.
- Log in using a certificate.
- Log in with SSO.

For information about managing user permissions, see the KB article *Data Domain - Managing User Permissions on the Data Domain system*, available from the Online Support website.

Related concepts

Managing host certificates for HTTP and HTTPS

Logging in with a user name and password

Log in to DD System Manager using a web browser and your assigned user name and password.

Steps

1. Open a web browser and enter the IP address or hostname to connect to DD System Manager. It must be one of the following:
   - A fully qualified domain name (for example, `http://dd01.example.com`)
   - A hostname (`http://dd01`)
   - An IP address (`http://10.5.50.5`)
2. For HTTPS secure login, click **Secure Login**.

   Secure login with HTTPS requires a digital certificate to validate the identity of the DDOS system and to support bi-directional encryption between DD System Manager and a browser. DDOS includes a self-signed certificate, and DDOS allows you to import your own certificate.
3. Enter your assigned username and password.
   - For physical systems the default password is the system serial number.
   - For PowerProtect DD Virtual Edition (DD VE) instances the default password is **changeme**.
4. Click **Log In**.

   On the first time login by the system administrator, the password must be changed from its default value.
   - On physical DD systems, the system administrator password must be changed from the system serial number to a new value that meets the password strength policies on the DD system.
   - On PowerProtect DD Virtual Edition (DDVE), the system administrator password must be changed from **changeme** to a new value that meets the password strength policies on the DDVE instance.

   For first-time login, the Home page appears.

   NOTE: If this is the first time you are logging in and the system administrator has configured your username to require a password change, you must change the password before gaining access to DD System Manager.
5. To log out, click the log out button in the DD System Manager banner.

Log in using CAC/PIV smart card and user certificates

Log in to DD System Manager with a certificate issued by a Certificate Authority (CA).

Prerequisites

- You must have authorization privileges on the protection system, and the protection system must trust the CA certificate. Your username must be specified in the common-name field in the certificate. For Active Directory users, specify the Microsoft UPN under OtherNames in the SubjectAlternativeName field instead of using the common-name field.
- You must have a user account on the protection system. You can be either a local user or a name service user (NIS/AD). For a name service user, your group-to-role mapping must be configured on the protection system.
- User certificates and upstream intermediate CAs should not be revoked by imported CRLs.

Steps

1. Use the following CLI command to import the public key from the CA that issued the certificate: `adminaccess certificate import ca application login-auth`.

   NOTE: If the CA certificate consists of a CA-chain, run the `adminaccess certificate import ca application login-auth` command multiple times to import each public key of the CA-chain up to the root CA.
2. Load the user certificate in PKCS12 format in your browser from the CAC/PIV card after swiping CAC/PIV card against a card reader which interacts with the browser .
   Once the CA certificate is trusted by the protection system, a **Log in with certificate** link is visible on the HTTPS login screen.
3. Click **Log in with certificate**, and choose the user certificate from the list of certificates that are prompted by the browser.

Results

The system validates the user certificate against the trust store. Based on authorization privileges associated with your account, a System Manager session is created for you.

Logging in using single sign-on (SSO)

Log in to DD System Manager with a username and password from a supported SSO provider.

Prerequisites

SSO must be enabled and the protection system must be registered with an SSO provider.

About this task

Configuring SSO authentication describes how to enable SSO authentication and register the protection system with the SSO provider.

Steps

1. At the login screen, click **Log in with Data Protection Central**.

   NOTE: If a brand name is set on Data Protection Central (DPC), the link appears as **Log in with *<DPC-brand-name>***.

2. Log in with the DPC username and password.

Logging in using multifactor authentication

Log in to DD System Manager with a DDSM username and password plus an RSA SecurID passcode.

Prerequisites

Multifactor authentication must be configured and enabled.

About this task

Configuring and enabling multifactor authentication (MFA) describes how to configure and enable multifactor authentication.

Steps

1. Enter your assigned username and password.
2. Click **Log In**.
3. Enter the RSA SecurID passcode.
4. Click **Verify**.

Troubleshooting login issues

The GUI Service is temporarily unavailable

DD System Manager is unable to launch from any web browser with this error message. Do one of the following:

- Refresh your browser. If the problem persists, contact support for assistance.
- Use SSH to login to the system and run all commands. If you are connecting through SSH, the SSH client must be running OpenSSH 4.7p1 or later.

  NOTE: If the system administrator logs in for the first time using SSH or Telnet, the system will force a password change from the default value. For fresh installed systems, the system also prompts you to configure a security officer. Select **yes** to provide a username and password for the security officer, or **no** to continue to the configuration wizard. For systems upgraded from an older version DDOS version, the system generates an alert about creating the security

officer but does not actively prompt for it. The *Managing System Access* chapter describes how to create a user with a security role.

- If you have not upgraded the DDOS and this GUI error occurs, use the following procedure:
  1. Close the web browser session on the system with the reported error.
  2. Run these commands in sequence:
     - `adminaccess disable http`
     - `adminaccess disable https`
     - `adminaccess enable http`
     - `adminaccess enable https`
  3. Wait 5 minutes to allow the http and https services to start completely.
  4. Open a web browser, and connect to DD System Manager.
- If this GUI issue occurs after a DDOS upgrade, use the following procedure:
  1. Close the web browser session on the system with the reported error.
  2. Run these commands in sequence:
     - `adminaccess disable http`
     - `adminaccess disable https`
     - `adminaccess certificate generate self-signed-cert`
     - `adminaccess enable http`
     - `adminaccess enable https`
  3. Wait 5 minutes to allow the http and https services to start completely.
  4. Open a web browser, and connect to DD System Manager.

User password is aged

When the user password status is `password-aged`, login with the current password to set a new password. If the aged password is for the security-officer, it cannot be used to authorize operations until a new password is set.

User locked out

If you enter an incorrect password 4 consecutive times, the system locks out the specified username for the duration specified in the `login-unlock-timeout` option (120 seconds is the default value). The login count and lockout period are configurable and might be different on your system.

Forgot sysadmin password or forgot the only security officer password

If you forget the sysadmin password or the only security officer password after changing it, contact Dell Support.

Using the system configuration wizard

The DD System Manager wizard guides you through a simplified configuration to get your system operating quickly. After you complete the basic configuration with a wizard, you can use DD System Manager and the CLI to further configure your system.

Steps

1. Select **Maintenance** > **System** > **Configure System**.
2. Use the controls at the bottom of the **Configuration Wizard** dialog box to select which features to configure and to advance through the wizard.

Configuration parameters

View the parameters that you can configure using the Configuration wizard.

The Configuration wizard enables you to define the parameters for the following components:

NOTE: For DD6400 systems, the Configuration Wizard displays different parameters.

- Licenses
- Network
- File System
- System settings
- Deployment Assessment page for DDVE
- DD Boost protocol
- CIFS protocol
- NFS protocol
- DD VTL protocol

The parameters for DD6400 systems are:

- License
- Network
- System Passphrase
- File System
- Settings

The Online Help provides more details about these options.

Using the command line interface

The command line interface (CLI) is a text-driven interface that you can use instead of or in addition to DD System Manager. Although most management tasks can be performed in DD System Manager, the CLI offers some configuration options and reports that are not yet supported in DD System Manager.

Any command that accepts a list, such as a list of IP addresses, accepts entries separated by commas, by spaces, or both.

The command-line interface is available through a serial console or through an Ethernet connection using SSH, Telnet, or serial over LAN (SOL). Some systems support access using a keyboard and monitor attached directly to the system.

The *DDOS Command Reference Guide* provides information for each of the CLI commands. Online help provides the complete syntax for each command.

Logging into the CLI

You can access the CLI by using a direct connection to the protection system or by using an Ethernet connection through SSH or Telnet. By default, SSH is enabled and Telnet is disabled.

Prerequisites

To use the CLI, you must establish a local or remote connection to the protection system using one of the following methods.

- If you are connecting through a serial console port on the system, connect a terminal console to the port and use the communication settings: 115200 baud, 8 data bits, no parity, and 1 stop bit.
- If the system supports keyboard and monitor ports, connect a keyboard and monitor to those ports.
- If you are connecting through Ethernet, connect a computer with SSH or Telnet client software to an Ethernet network that can communicate with the system.
- If you are connecting through SSH, the SSH client must be running OpenSSH 4.7p1 or later.

Steps

1. If you are using an SSH or Telnet connection to access the CLI, start the SSH or Telnet client and specify the IP address or host name of the protection system.

   For information on initiating the connection, see the documentation for the client software. The system prompts you for a username.
2. When prompted, enter your protection system username or **`sysadmin`**, the default username.
3. When prompted, enter the password for the specified username.

Example

The following example shows SSH login to a system named *mysystem* using SSH client software.

```
# ssh -l sysadmin mysystem.mydomain.com
 DD9900-157.dellemc.com
DDOS
Password:
```

CLI online help guidelines

The CLI displays two types of help: syntax-only help and command-description help, which includes the command syntax. Both types of help offer features that enable you reduce the time it takes to find the information you need.

The following guidelines describe how to use syntax-only help.

- To list the top-level CLI commands, enter a question mark (**?**), or type `help` or `man` at the prompt.
- To list all forms of a top-level command, enter the command with no options at the prompt or enter *command* `?`.
- To list all commands that use a specific keyword, enter `help` *keyword*, `man` *keyword*, or `?` *keyword*.

  For example, `? password` displays all system commands that use the password argument.

The following guidelines describe how to use command-description help.

- To list the top-level CLI commands, enter a question mark (**?**), or type `help` or `man` at the prompt.
- To list all forms of a top-level command with an introduction, enter `help` *command*, `man` *command*, or `?` *command*.
- The end of each help description is marked `END`. Press Enter to return to the CLI prompt.
- When the complete help description does not fit in the display, the colon prompt (:) appears at the bottom of the display. The following guidelines describe what you can do when this prompt appears.
  - To move through the help display, use the up and down arrow keys.
  - To quit the current help display and return to the CLI prompt, press **q**.
  - To display help for navigating the help display, press **h**.
  - To search for text in the help display, enter a slash character (**/**) followed by a pattern to use as search criteria and press Enter. Matches are highlighted.

Managing HA systems

The High Availability (HA) feature lets you configure two protection systems as an Active-Standby pair, providing redundancy in the event of a system failure. HA keeps the active and standby systems in sync, so that if the active node were to fail due to hardware or software issues, the standby node can take over services and continue where the failing node left off.

Using DD System Manager, you can view the status of a configured HA system, but you cannot perform an initial HA system set-up. Use the DDOS CLI commands to set up the HA relationship between the two nodes, one active and one standby.

HA is supported on the following systems:

- Data Domain DD6800
- Power Protect DD6900
- Data Domain DD9300
- Power Protect DD9400
- Data Domain DD9800
- Power Protect DD9900

Setup

Both nodes of the HA pair must have identical hardware. This requirement is validated during setup and system boot-up. Ensure that the system interconnect and identical hardware are set up on both nodes. Run the initial set-up on either node, one at a time.

When configuring HA for the first time, run the `ha create` command on the node with the license installed. To upgrade an existing system to HA by adding a new or unconfigured system, initiate the HA upgrade from the existing standalone system.

Feature continuity

HA provides failover within 10 minutes for most operations. CIFS, DD VTL, and NDMP must be restarted manually.

NOTE: Recovery of DD Boost applications may take longer than 10 minutes, because Boost application recovery cannot begin until the DD server failover is complete. In addition, Boost application recovery cannot start until the application invokes the Boost library. Similarly, NFS may require additional time to recover.

The installation guides for the systems that support HA describe how to install a new HA system. The *Single Node to HA Upgrade* describes how to upgrade an existing system to an HA pair.

Maintenance

The HA architecture provides a rolling upgrade, which reduces maintenance downtime for the upgrade.

With a rolling upgrade, the HA nodes are upgraded one at a time. The standby node is restarted and upgraded first. The newly upgraded standby node then takes over the active role through an HA failover. After the failover, the second node is restarted and assumes the role of the standby node after the upgrade.

System upgrade operations that require data conversion cannot start until both systems are upgraded to the same level and HA state is fully restored.

Managing electronic licenses

Add and delete electronic licenses from the system. Refer to the applicable Release Notes for the most up-to-date information on product features, software updates, software compatibility guides, and information about products, licensing, and service. Use the DDSM GUI or the DDSH CLI to add, replace, or delete electronic licenses. When using the GUI to manage licenses, the system accepts license files with `.lic`, `.txt`, or `.xml` file extensions. However, if DD Retention Lock Compliance is enabled on the system, the `elicense` CLI commands are required to update or remove licenses because license control in the GUI is disabled.

Managing HA system licenses

HA is a licensed feature, and the system licensing key is registered by following the steps to add any other license to the DD system.

A system is configured as Active-Standby, where one node is designated "standby." Only one set of licenses is required for both nodes. During failover, the licenses on one node will failover to the other node.

Optionally configure the login banner

DD System Manager provides the ability to display a custom message on the login screen to inform users that the system contains confidential data, and only authorized personnel are allowed access.

Steps

1. Select **Administration** > **Access**.
2. Select the **Login Banner** tab.
3. Click **Configure** to configure the login banner for the first time.
4. In the Title field, specify a title for the login banner message.
5. Select **Message text** to type a message, or **Upload a .txt file** to upload a message from a text file.
6. Click **Save**.

Modifying or resetting the login banner

About this task

The system administrator can modify or reset the login banner message at any time after it is configured.

Steps

1. Click **Modify** to change the text of the message.
2. Click **Reset** to remove the message entirely.

3

Configuring System Settings

This chapter presents the following topics:

Topics:

Managing the system passphrase

The system passphrase is a key that allows a protection system to be transported with encryption keys on the system. The encryption keys protect the data and the system passphrase protects the encryption keys.

The system passphrase is a password-like phrase, which generates an AES 256 encryption key. If the system is stolen in transit, an attacker cannot easily recover the data; at most, they can recover the encrypted user data and the encrypted keys.

The passphrase is stored internally on a hidden part of the storage subsystem. This allows the protection system to boot and continue servicing data access without administrator intervention.

Setting the system passphrase

The system passphrase must be set before the system can support data encryption or request digital certificates.

Prerequisites

No minimum system passphrase length is configured when DDOS is installed, but the CLI provides a command to set a minimum length. To determine if a minimum length is configured for the passphrase, enter the `system passphrase option show` CLI command.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.

   If the system passphrase is not set, the **Set Passphrase** button appears in the Passphrase area. If a system passphrase is configured, the **Change Passphrase** button appears, and your only option is to change the passphrase.

2. Click the **Set Passphrase** button.

   The Set Passphrase dialog appears.

3. Enter the system passphrase in the boxes and click **Next.**

   If a minimum length is configured for the system passphrase, the passphrase you enter must contain the minimum number of characters.

Results

The system passphrase is set and the **Change Passphrase** button replaces the **Set Passphrase** button.

Changing the system passphrase

The administrator can change the passphrase without having to manipulate the encryption keys. Changing the passphrase indirectly changes the encryption of the keys, but does not affect user data or the underlying encryption key.

About this task

WARNING: Be sure to take care of the passphrase. If the passphrase is lost, it cannot be recovered.

Changing the passphrase requires two-user authentication to protect against data shredding.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. To change the system passphrase, click **Change Passphrase**.

   The Change Passphrase dialog appears.

   NOTE: The file system must be disabled to change the passphrase. If the file system is running, you are prompted to disable it.

3. In the text fields, provide:
   - The username and password of a Security Officer account (an authorized user in the Security User group on that system).
   - The current passphrase when changing the passphrase.
   - The new passphrase, which must contain the minimum number of characters that are configured with the `system passphrase option set min-length` command.
4. Click the checkbox for **Enable file system now**.
5. Click **OK**.

Enabling FIPS mode

The FIPS mode button allows you to enable or disable FIPS 140-2 compliance mode.

Prerequisites

The *DDOS, PowerProtect DD Virtual Edition, and PowerProtect DD Management Center Security Configuration Guide* provides additional details about FIPS 140-2 compliance on DDOS.

Steps

1. Select **Administration** > **Settings**.
2. Click **FIPS Mode** to enable or disable FIPS 140-2 compliance mode.

Results

After enabling FIPS 140-2 compliance mode, DDOS:

- Forces a password change for the sysadmin account and one security officer account (if security officer is enabled).
- Reboots, causing an interruption in file system access.
- Allows only applications with FIPS-compatible clients to access the file system after the reboot is complete.

Configuring mail server settings

The Mail Server tab allows you to specify the mail server to which DDOS sends email reports.

About this task

Steps

1. Select **Administration** > **Settings** > **Mail Server**.
2. Select **More Tasks** > **Set Mail Server**.

   The Set Mail Server dialog box appears.
3. Specify the name of the mail server in the **Mail Server** field.
4. Use the **Credentials** button to enable or disable the use of credentials for the mail server.
5. If credentials are enabled, specify the mail server username in the **User Name** field.
6. If credentials are enabled, specify the mail server password in the **Password** field.
7. Click **Set**.
8. Optionally use the CLI to verify and troubleshoot the mail server configuration.
   a. Run the **`config show mailserver`** command to verify the mail server is configured.
   b. Run the **`net ping <mailserver-hostname> count 4`** command to ping the mail server.
   c. If the mail server is not configured correctly, run the **`config set mailserver <mailserver-hostname>`** command to set the mail server, and attempt to ping it again.
   d. Run the **`net show dns`** command to verify the DNS server is configured.
   e. Run the **`net ping <DNS-hostname> count 4`** command to ping the DNS server.
   f. If the DNS server is not configured correctly, run the **`config set dns <dns-IP>`** command to set the DNS server, and attempt to ping it again.
   g. Optionally run the **`net hosts add <IP-address> <hostname>`** command to add the mail server IP address and hostname to the system hosts file for local resolving.
   h. Run the **`net ping <mailserver-hostname> count 4`** command to ping the mail server.

Managing time and date settings

The Time and Date Settings tab allows you to view and configure the system time and date or configure the Network Time Protocol to set the time and date.

Steps

1. To view the current time and date configuration, select **Administration** > **Settings** > **Time and Date Settings**.

   The Time and Date Settings page presents the current system date and time, shows whether NTP is enabled or not, and lists the IP addresses or hostnames of configured NTP servers.
2. To change the configuration, select **More Tasks** > **Configure Time Settings**.

   The Configure Time Settings dialog appears.
3. In the **Time Zone** dropdown list, select the time zone where the Data Domain system resides.
4. To manually set the time and date, select **None**, type the date in the **Date** box, and select the time in the **Time** dropdown lists.
5. To use NTP to synchronize the time, select NTP and set how the NTP server is accessed.
   - To use DHCP to automatically select a server, select **Obtain NTP Servers using DHCP**.
   - To configure an NTP server IP address, select **Manually Configure**, add the IP address of the server, and click **OK**.

   NOTE: Using time synchronization from an Active Directory domain controller might cause excessive time changes on the system if both NTP and the domain controller are modifying the time.
6. Click **OK**.
7. If you changed the time zone, you must reboot the system.
   a. Select **Maintenance** > **System**.
   b. From the More Tasks menu, select Reboot System.
   c. Click OK to confirm.

Configuring secure NTP with symmetric keys

Configure secure NTP from the CLI.

Prerequisites

Configure secure NTP settings on the time server before adding the time server to the protection system. When adding the time server to the protection system, the system validates the key ID and key value on the time server.

Steps

1. Run the **ntp secure add authentication-key *<key-ID>* SHA1** command to add the authentication key ID and key value.

```
# ntp secure add authentication-key 1 SHA1
Enter the NTP authentication key:
Re-enter the NTP authentication key:
Key matched.
The authentication key is added successfully.
```

2. Run the **ntp secure add timeserver *<timeserver>* *<key-ID>*** command to add the time server and corresponding authentication key.

   NOTE: The time server is added if the specified key ID and key value are an exact match to the key ID and key value configured on the time server.

```
# ntp secure add timeserver A.B.C.D 1
NTP server is added successfully
Remote Time Servers    Key Number
-------------------    ----------
A.B.C.D                1
-------------------    ----------
```

3. Run the **ntp secure add trusted-key *<key-ID>*** command to add the authentication key to the trusted key list.

```
# ntp secure add trusted-key 1
The authentication key ID is added to trusted key list successfully.
```

4. Run the **ntp secure show config** command to check the configuration.

```
# ntp secure show config
Secure NTP is currently disabled.
#    Servers           Key Number
-    --------------    ----------
1    A.B.C.D           1
-    --------------    ----------

Authentication Key ID's : 1, 2, 3

Trusted Key ID's : 1

Showing NTP servers configured manually.
```

5. Run the **ntp secure status** command to check the secure NTP status.

```
# ntp secure status
Status                                      Disabled
Current Clock Time                          Unknown
Clock Last Synchronized                     Unknown
Clock Last Synchronized With Time Server    Unknown
```

6. Run the **ntp secure enable** command to enable secure NTP.

```
# ntp secure enable
secure NTP enabled.
```

7. Run the **ntp secure sync** command to synchronize the system date and time with the secure NTP time server.

```
# ntp secure sync
****  Current time is set to Thu Dec 8 18:33:34.461 2022.
      It may take a few minutes until the system time is fully synchronized.
```

Setting system date change frequency and date change limit

Optionally configure the system-enforced interval between system time and date changes, and the maximum allowed amount to advance the system time and date. These values only take effect after DD Retention Lock Compliance is enabled and cannot be configured afterwards.

About this task

This task must be performed through the CLI. Full administrative privileges with security officer oversight are required. Limited-admin users are not permitted to run these commands.

Steps

1. Run the `system set date-change-frequency` command to set the allowed interval between system time and date changes.

   ```
   #  system set date-change-frequency [<DD> | reset]
   ```

   Where ***<DD>*** is the number of days required between time and date changes.

2. Run the `system set date-change-limit` command to set the maximum allowed advance for a single system time and date change operation.

   ```
   #  system set set date-change-limit [<hh:mm> | reset]
   ```

   Where ***<hh>*** is the number of hours, and ***<mm>*** is the number of minutes.

   Once the date change limit is set, the system generates an alert when the clock skew exceeds half of the date change limit. If the alert appears, fix the system time and clear the alert manually. If the alert is not cleared, it will update for any further increase in the clock skew (when the clock skew increases by at least half of the system date change limit).

Managing system properties

The System Properties tab enables you to view and configure system properties that identify the managed system location, administrator email address, and host name.

Steps

1. To view the current configuration, select **Administration** > **Settings** > **System Properties**.

   The System Properties tab displays the system location, the administrator email address, and the administrator hostname.

2. To change the configuration, select **More Tasks** > **Set System Properties**.

   The Set System Properties dialog box appears.

3. In the **Location** box, enter information about where the protection system is located.
4. In the **Admin Email** box, enter the email address of the system administrator.
5. In the **Admin Host** box, enter the name of the administration server.
6. Click **OK**.

Managing SNMP

The Simple Network Management Protocol (SNMP) is a standard protocol for exchanging network management information, and is a part of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite. SNMP provides a tool for network administrators to manage and monitor network-attached devices, such as DD systems, for conditions that warrant administrator attention.

To monitor systems using SNMP, install the DDOS MIB in your SNMP Management system. DDOS also supports the standard MIB-II so you can query MIB-II statistics for general data such as network statistics. For full coverage of available data, use both the DDOS MIB and the standard MIB-II.

The DDOS system SNMP agent accepts queries for system-specific information from management systems using SNMP v1, v2c, and v3. SNMP V3 provides a greater degree of security than v2c and v1 by replacing cleartext community strings (used for authentication) with user-based authentication using either MD5, SHA1, or SHA256. SNMP v3 user authentication packets can be encrypted and their integrity verified with either DES or AES.

DD systems can send SNMP traps (which are alert messages) using SNMP v2c and SNMP v3. Because SNMP v1 traps are not supported, if possible, use SNMP v2c or v3.

The default port that is open when SNMP is enabled is port 161. Traps are sent out through port 162.

The *DDOS MIB Quick Reference* describes the full set of MIB parameters included in the DDOS MIB branch.

Viewing SNMP status and configuration

The SNMP tab displays the current SNMP status and configuration.

Steps

Select **Administration** > **Settings** > **SNMP**.

The SNMP view shows the SNMP status, SNMP properties, SNMP V3 configuration, and SNMP V2C configuration.

Enabling and disabling SNMP

Use the SNMP tab to enable of disable SNMP.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the Status area, click **Enable** or **Disable**.

Downloading the SNMP MIB

Use the SNMP tab to download the SNMP MIB.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. Click **Download MIB file**.
3. In the Opening *<protection system>*.mib dialog box, select **Open**.
4. Click **Browse** and select a browser to view the MIB in a browser window.

   NOTE: If using the Microsoft Internet Explorer browser, enable Automatic prompting for file download.

5. Save the MIB or exit the browser.

Configuring SNMP properties

Use the SNMP tab to configure the text entries for system location and system contact.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the SNMP Properties area, click **Configure**.

   The SNMP Configuration dialog box appears.

3. In the text fields, specify the following information: and/or an
   - SNMP System Location: A description of where the protection system is located.

NOTE: For HA pairs, this value can also indicate whether the system is node 0 or node 1.

- SNMP System Contact: The email address of the system administrator.
- SNMP System Notes: (Optional) Additional SNMP configuration information.
- SNMP Engine ID: A unique identifier for the SNMP entity. The following requirements and guidelines apply:
  - The engine ID must be 34 hexadecimal characters (SNMPv3 only).
    NOTE: The system displays an error if the SNMP engine ID does not meet the length requirements, or uses invalid characters.
  - Use a value meaningful to the installation.
  - For HA pairs, the engine ID can only be changed from the active node, and will be the same for all nodes.

4. Click **OK**.

SNMP V3 user management

Use the SNMP tab to create, modify, and delete SNMPv3 users and trap hosts.

Creating SNMP V3 users

When you create SNMPv3 users, you define a username, specify either read-only or read-write access, and select an authentication protocol.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the SNMP Users area, click **Create**.

   The Create SNMP User dialog box appears.

3. In the **Name** text field, enter the name of the user for whom you want to grant access to the system agent. The name must be a minimum of eight characters.
4. Select either read-only or read-write access for this user.
5. To authenticate the user, select **Authentication**.
   a. Select the MD5, SHA1, or SHA256 protocol.
   b. Enter the authentication key in the **Key** text field.
   c. To provide encryption to the authentication session, select **Privacy**.
   d. Select either the AES or the DES protocol.
   e. Enter the encryption key in the **Key** text field.
6. Click **OK**.

   The newly added user account appears in the SNMP Users table.

Modifying SNMP V3 users

You can modify the access level (read-only or read-write) and authentication protocol for existing SNMPv3 users.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the **SNMP Users** area, select a checkbox for the user and click **Modify**.

   The Modify SNMP User dialog box appears. Add or change any of the following settings.

3. Select either read-only or read-write access for this user.
4. To authenticate the user, select **Authentication**.
   a. Select the MD5, SHA1, or SHA256 protocol.
   b. Enter the authentication key in the Key text field.
   c. To provide encryption to the authentication session, select **Privacy**.
   d. Select either the AES or the DES protocol.

e. Enter the encryption key in the **Key** text field.

5. Click **OK**.

   The new settings for this user account appear in the SNMP Users table.

Removing SNMP V3 users

Use the SNMP tab to delete existing SNMPv3 users.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the SNMP Users area, select a checkbox for the user and click **Delete**.

   The Delete SNMP User dialog box appears.

   NOTE: If the **Delete** button is disabled, the selected user is being used by one or more trap hosts. Delete the trap hosts and then delete the user.

3. Verify the user name to be deleted and click **OK**.
4. In the Delete SNMP User Status dialog box, click **Close**.

   The user account is removed from the SNMP Users table.

SNMP V2C community management

Define SNMP v2c communities (which serve as passwords) to control management system access to the protection system. To restrict access to specific hosts that use the specified community, assign the hosts to the community.

NOTE: The SNMP V2c Community string is a sent in cleartext and is very easy to intercept. If this occurs, the interceptor can retrieve information from devices on your network, modify their configuration, and possibly shut them down. SNMP V3 provides authentication and encryption features to prevent interception.

NOTE: SNMP community definitions do not enable the transmission of SNMP traps to a management station. You must define trap hosts to enable trap submission to management stations.

Creating SNMP V2C communities

Create communities to restrict access to the DDR system or for use in sending traps to a trap host. You must create a community and assign it to a host before you can select that community for use with the trap host.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the Communities area, click **Create**.

   The Create SNMP V2C Community dialog box appears.

3. In the **Community** box, enter the name of a community for whom you want to grant access to the system agent.
4. Select either read-only or read-write access for this community.
5. If you want to associate the community to one or more hosts, add the hosts as follows:
   a. Click **+** to add a host.

      The Host dialog box appears.

   b. In the **Host** text field, enter the IP address or domain name of the host.
   c. Click **OK**.

      The Host is added to the host list.

6. Click **OK**.

   The new community entry appears in the **Communities** table and lists the selected hosts.

Modifying SNMP V2C Communities

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the Communities area, select the checkbox for the community and click **Modify**.

   The Modify SNMP V2C Community dialog box appears.

3. To change the access mode for this community, select either **read-only** or **read-write** access.

   NOTE: The Access buttons for the selected community are disabled when a trap host on the same system is configured as part of that community. To modify the access setting, delete the trap host and add it back after the community is modified.

4. To add one or more hosts to this community, do the following:
   a. Click **+** to add a host.

      The Host dialog box appears.
   b. In the **Host** text field, enter the IP address or domain name of the host.
   c. Click **OK**.

      The Host is added to the host list.
5. To delete one or more hosts from the host list, do the following:

   NOTE: DD System Manager does not allow you to delete a host when a trap host on the same system is configured as part of that community. To delete a trap host from a community, delete the trap host and add it back after the community is modified.

   NOTE: The Access buttons for the selected community are not disabled when the trap host uses an IPv6 address and the system is managed by an earlier DDOS version that does not support IPv6. If possible, always select a management system that uses the same or a newer DDOS version than the systems it manages.

   a. Select the checkbox for each host or click the Host check box in the table head to select all listed hosts.
   b. Click the delete button (X).
6. To edit a host name, do the following:
   a. Select the checkbox for the host.
   b. Click the edit button (pencil icon).
   c. Edit the host name.
   d. Click **OK**.
7. Click **OK**.

   The modified community entry appears in the Communities table.

Deleting SNMP V2C communities

Use the SNMP tab to delete existing SNMPv2 communities.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the **Communities** area, select a checkbox for the community and click **Delete**.

   The Delete SNMP V2C Communities dialog box appears.

   NOTE: If the **Delete** button is disabled, the selected community is being used by one or more trap hosts. Delete the trap hosts and then delete the community.

3. Verify the community name to be deleted and click **OK**.
4. In the Delete SNMP V2C Communities Status dialog box, click **Close**. The community entry is removed from the Communities table.

SNMP trap host management

Trap host definitions enable protection systems to send alert messages in SNMP trap messages to an SNMP management station.

Creating SNMP V3 and V2C trap hosts

Trap host definitions identify remote hosts that receive SNMP trap messages from the system.

Prerequisites

If you plan to assign an existing SNMP v2c community to a trap host, you must first use the Communities area to assign the trap host to the community.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the SNMP V3 Trap Hosts or SNMP V2C Trap Hosts area, click **Create**.

   The Create SNMP [V3 or V2C] Trap Hosts dialog appears.
3. In the **Host** box, enter the IP address or domain name of the SNMP Host to receive traps.
4. In the **Port** box, enter the port number for sending traps (port 162 is a common port).
5. Select the user (SNMP V3) or the community (SNMP V2C) from the drop-down menu.

   NOTE: The Community list displays only those communities to which the trap host is already assigned.
6. To create a new community, do the following:
   a. Select **Create New Community** in the Community drop-down menu.
   b. Enter the name for the new community in the **Community** box.
   c. Select the Access type.
   d. Click the add (+) button.
   e. Enter the trap host name.
   f. Click **OK.**
   g. Click **OK.**
7. Click **OK**.

Modifying SNMP V3 and V2C trap hosts

You can modify the port number and community selection for existing trap host configurations.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the **SNMP V3 Trap Hosts** or **SNMP V2C Trap Hosts** area, select a Trap Host entry, and click **Modify**.

   The Modify SNMP [V3 or V2C] Trap Hosts dialog box appears.
3. To modify the port number, enter a new port number in the **Port** box (port 162 is a common port).
4. Select the user (SNMP V3) or the community (SNMP V2C) from the drop-down menu.

   NOTE: The Community list displays only those communities to which the trap host is already assigned.
5. To create a new community, do the following:
   a. Select **Create New Community** in the Community drop-down menu.
   b. Enter the name for the new community in the **Community** box.
   c. Select the Access type.
   d. Click the add (+) button.
   e. Enter the trap host name.
   f. Click **OK.**

g. Click **OK.**

6. Click **OK**.

Removing SNMP V3 and V2C trap hosts

Use the SNMP tab to delete existing trap host configurations.

Steps

1. Select **Administration** > **Settings** > **SNMP**.
2. In the **Trap Hosts** area (either for V3 or V2C, select a checkbox for the trap host and click **Delete**.

   The Delete SNMP [V3 or V2C] Trap Hosts dialog box appears.

3. Verify the host name to be deleted and click **OK**.
4. In the Delete SNMP [V3 or V2C] Trap Hosts Status dialog box, click **Close**.

   The trap host entry is removed from the **Trap Hosts** table.

Configuring the cipher-list for SSL and TLS

SMS and the DDSM GUI support all the ciphers listed in the table below, but because of the default configuration for the SMS context and DDSM context, only a small number of cipher-lists are supported by each service. The default cipher-list is `DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-RSA-AES128-GCM-SHA256`, as shown by the `adminaccess option show cipher-list` command.

To support backwards compatibility, use the **`adminaccess option set cipher-list <cipher-list>`** command to add ciphers from the following table.

Supported cipher-lists
ECDHE-RSA-AES256-GCM-SHA384
ECDHE-ECDSA-AES256-GCM-SHA384
ECDHE-RSA-AES256-SHA384
ECDHE-ECDSA-AES256-SHA384
ECDHE-RSA-AES256-SHA
ECDHE-ECDSA-AES256-SHA
SRP-DSS-AES-256-CBC-SHA
SRP-RSA-AES-256-CBC-SHA
SRP-AES-256-CBC-SHA
DHE-DSS-AES256-GCM-SHA384
DHE-RSA-AES256-GCM-SHA384
DHE-RSA-AES256-SHA256
DHE-DSS-AES256-SHA256
DHE-RSA-AES256-SHA
DHE-DSS-AES256-SHA
DHE-RSA-CAMELLIA256-SHA
DHE-DSS-CAMELLIA256-SHA
AECDH-AES256-SHA
ECDH-RSA-AES256-GCM-SHA384
ECDH-ECDSA-AES256-GCM-SHA384
ECDH-RSA-AES256-SHA384

Supported cipher-lists
ECDH-ECDSA-AES256-SHA384
ECDH-RSA-AES256-SHA
ECDH-ECDSA-AES256-SHA
AES256-GCM-SHA384
AES256-SHA256
AES256-SHA
CAMELLIA256-SHA
PSK-AES256-CBC-SHA
ECDHE-RSA-AES128-GCM-SHA256
ECDHE-ECDSA-AES128-GCM-SHA256
ECDHE-RSA-AES128-SHA256
ECDHE-ECDSA-AES128-SHA256
ECDHE-RSA-AES128-SHA
ECDHE-ECDSA-AES128-SHA
SRP-DSS-AES-128-CBC-SHA
SRP-RSA-AES-128-CBC-SHA
SRP-AES-128-CBC-SHA
DHE-DSS-AES128-GCM-SHA256
DHE-RSA-AES128-GCM-SHA256
DHE-RSA-AES128-SHA256
DHE-DSS-AES128-SHA256
DHE-RSA-AES128-SHA
DHE-DSS-AES128-SHA
DHE-RSA-CAMELLIA128-SHA
DHE-DSS-CAMELLIA128-SHA
AECDH-AES128-SHA
ECDH-RSA-AES128-GCM-SHA256
ECDH-ECDSA-AES128-GCM-SHA256
ECDH-RSA-AES128-SHA256
ECDH-ECDSA-AES128-SHA256
ECDH-RSA-AES128-SHA
ECDH-ECDSA-AES128-SHA
AES128-GCM-SHA256
AES128-SHA256
AES128-SHA
CAMELLIA128-SHA
PSK-AES128-CBC-SHA
ECDHE-RSA-NULL-SHA
ECDHE-ECDSA-NULL-SHA
AECDH-NULL-SHA
ECDH-RSA-NULL-SHA
ECDH-ECDSA-NULL-SHA

Supported cipher-lists
NULL-SHA NULL-MD5

Troubleshooting system management

Slow response time

When processing a heavy load, a system might be less responsive than normal. In this case, management commands issued from either DD System Manager or the CLI might take longer to complete. When the duration exceeds allowed limits, a timeout error is returned, even if the operation completed.

The following table lists the recommendations for the maximum number of user sessions supported by DD System Manager:

Table 3. Maximum number of users supported by DD System Manager

System Model	Maximum Active Users	Maximum Logged In Users
16 GB and greater models[a]	10	20

a. DD6300, DD6800, DD6900, DD9300, DD9400, DD9800, and DD9900

4

Monitoring DD Systems

This chapter presents the following topics:

Topics:

Monitoring system information

The **Dashboard** displays summary information and status for alerts, the file system, licensed services, and hardware enclosures. The **Maintenance** area displays additional system information, including the system uptime and system and chassis serial numbers.

Steps

- To view system dashboard, select **Home** > **Dashboard**.

  From the Dashboard you can view the following information:
  - Alerts—Shows the most recent alerts for each subsystem (hardware, replication, file system, and others).
  - File System—Shows statistics for the entire file system.
  - Dashboard services—Shows the status of replication, DD VTL, CIFS, NFS, DD Boost, and vDisk services.
  - HA Readiness—Indicates whether the system can fail over from the active node to the standby node if necessary.
  - Hardware—Shows the status of the system enclosures and drives.
  - Maintenance—Shows the system model number, DDOS version, system uptime, and system and chassis serial numbers.
- To view the system uptime and identity information, select **Maintenance** > **System**.

Viewing hardware component status

The Hardware Chassis panel displays a block drawing of each enclosure in a system, including the chassis serial number and the enclosure status. Within each block drawing are the enclosure components, such as disks, fans, power supplies, NVRAM, CPUs, and memory. The components that appear depend upon the system model.

About this task

DD SM also displays the system serial number. The system serial number is independent of the chassis serial number and remains the same during many types of maintenance events, including chassis replacements.

Steps

1. Select **Hardware** > **Chassis**.

   The Chassis view shows the system enclosures. Enclosure 1 is the system controller, and the rest of the enclosures appear below Enclosure 1.

   Components with problems show yellow (warning) or red (error); otherwise, the component displays OK.

2. Click a component to see detailed status.

Viewing system statistics

The Realtime Charts panel displays up to seven charts that show real-time subsystem performance statistics, such as CPU usage and disk traffic.

Steps

1. Select **Home** > **Realtime Charts**.
   The Performance Charts area displays the currently selected charts.
2. To view specific data-point information, hover over a chart point.
3. When a chart contains multiple data, you can use the checkboxes in the upper-right corner of the chart to select what to display. For example, if Read is not selected in the upper right of the disk activity chart, only write data is charted.

Results

Each chart shows usage over the last 5 to 10 minutes.

Capacity statistics charts

DD System Manager enables you to view statistics for the amount of data backup to the DD system, and the amount of deduplication performed on the system in chart form.

Navigate to **Data Management** > **File System** > **Charts** to view the charts. There are three different charts available:

- Space Usage
- Consumption
- Daily Written

Use the **Chart** and **Scope** list boxes to select the parameters for the chart to display.

NOTE: For systems without Cloud Tier, **File System** is the only available scope. For systems with Cloud Tier, the scope can be **File System**, **Active Tier**, or **Cloud Tier**.

Use the **Date Range** options to select a date range to display on the chart. Choose from one of the preconfigured time periods, or specify specific dates to view.

Customize any chart by selecting or deselecting the available labels.

Space Usage chart

This chart displays a cumulative representation of how much data is backed up to the system, the amount of deduplication performed on that data, and the amount of space consumed on the system.

The space usage chart uses the following labels:

- Pre-Comp Used: The amount of storage consumed before compression.
- Post-Comp Used: The amount of storage consumed after compression.
  NOTE: This value is based on historical statistics and does not take into account any deletions. Use PCM to find a current estimate for post-compression storage usage.
- Comp Factor: The compression ratio.

Consumption chart

This chart displays a cumulative representation of the amount of space consumed on the system, and the amount of deduplication performed against the total capacity of the system. Administrators may also view the time lines and durations for system cleaning and data movement operations.

The space usage chart uses the following labels:

- Capacity: The total capacity of the system.
- Post-Comp Used: The amount of storage consumed after compression.
- Comp Factor: The compression ratio.

- Cleaning: The time lines and duration of system cleaning operations.
- Data Movement: The time lines and duration of data movement operations.

Daily Written chart

This chart displays the amount of data written to the system and the amount of deduplication performed on the system on a daily basis.

The space usage chart uses the following labels:

- Pre-Comp Written: The amount of data written to the system before compression.
- Post-Comp Written: The amount of data written to the system after compression.
- Total Comp Factor: The total compression ratio.

Viewing the Task Log

The Task Log displays a list of currently running jobs, such as, replication or system upgrades. DD System Manager can manage multiple systems and can initiate tasks on those systems. If a task is initiated on a remote system, the progress of that task is tracked in the management station task log, not in the remote system task log.

Steps

1. Select **Health** > **Jobs**.

   The Tasks view appears.

2. Select a filter by which to display the Task Log from the Filter By list box. You can select **All**, **In Progress**, **Failed**, or **Completed**.

   The Tasks view displays the status of all tasks based on the filter you select and refreshes every 60 seconds.

3. To manually refresh the Tasks list, do either of the following.
   - Click **Update** to update the task log.
   - Click **Reset** to display all tasks and remove any filters that were set.
4. To display detailed information about a task, select the task in the task list.

Viewing the system High Availability status

You can use the **High Availability** panel to see detailed information about the HA status of the system and whether the system can perform failover if necessary.

Steps

1. Select **Health > High Availability** on the DD System Manager.
   The **Health High Availability** screen appears.

   A green check mark indicates the system is operating normally and ready for failover.

   The screen shows the active node, which is typically Node 0.
2. Hover the cursor over a node to see its status.

   The node is highlighted in blue if it is active.
3. Click the drop-down menu in the banner if you want to change the view from the active node to the standby node, which is typically Node 1.

High Availability status

The **Health High Availability** (HA) view informs you about the system status using a diagram of the nodes and their connected storage. You can also see any current alerts as well as detailed information about the system, including the HA failover history.

You can determine if the active node and the storage are operational by hovering the cursor over them. Blue highlighting indicates normal operation. The standby node should appear gray.

You can filter the alerts table by selecting a component. Only alerts related to the selected components are displayed.

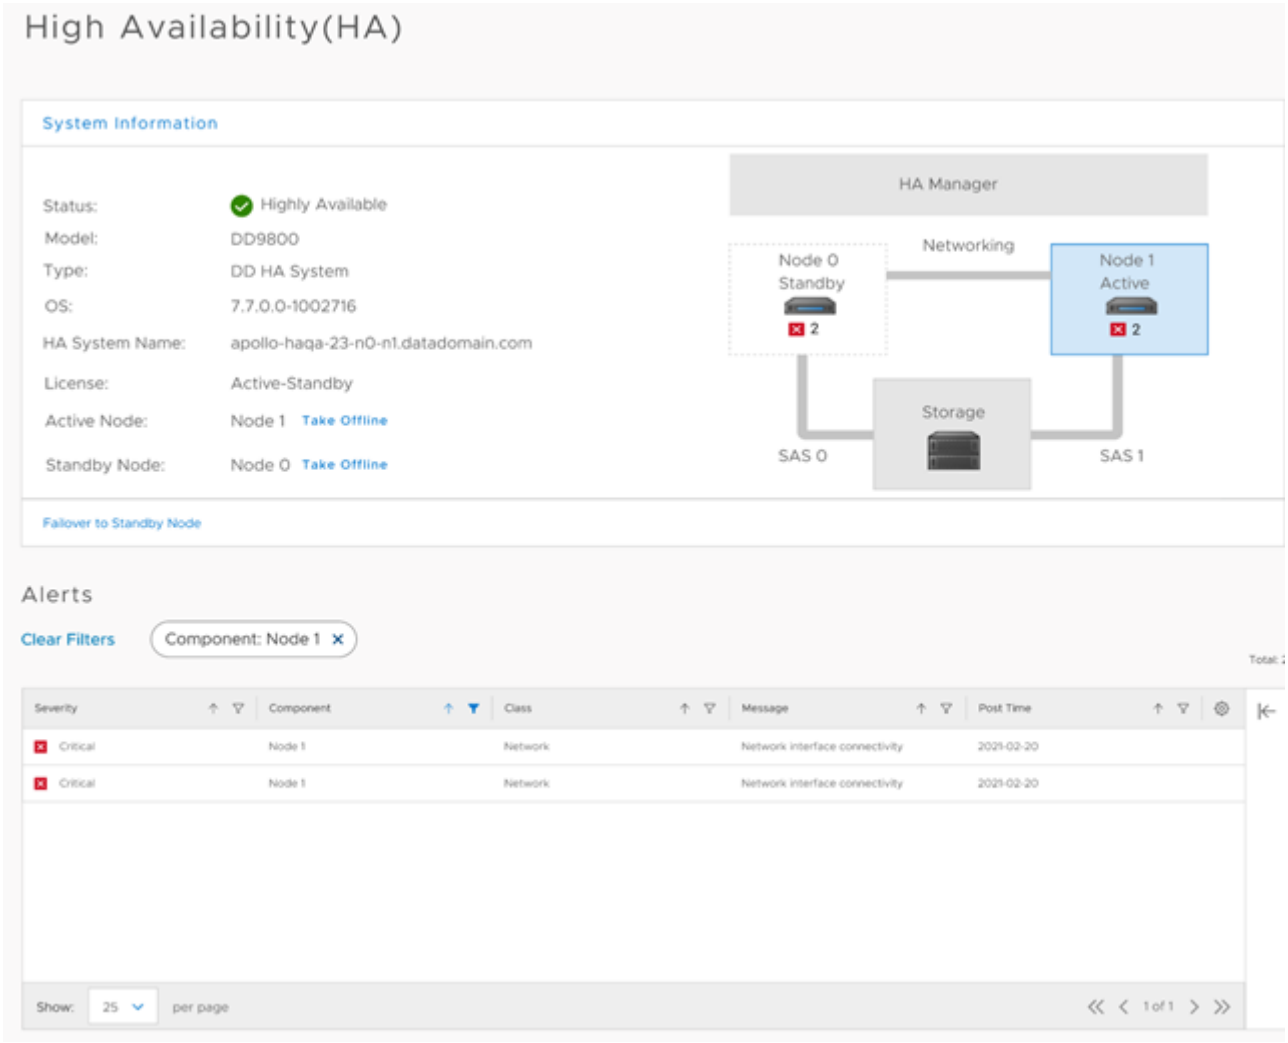

Figure 1. Health/High Availability indicators

5

Managing System Power

This chapter presents the following topics:

Topics:

Restarting a DDOS system

After modifying the system configuration, you might need to restart the system for the change to take effect. For example, changing the time zone requires that you restart the system before the new time zone is applied.

Steps

1. Select **Maintenance** > **System** > **Reboot System**.
2. Click **OK** to confirm.

Powering the DDOS system off

When powering a protection system off, follow the proper procedure to preserve the file system and configuration integrity.

About this task

CAUTION: Do not use the chassis power switch or the IPMI Remote System Power Down feature to power off the protection system unless the `system poweroff` command is unsuccessful. Using the chassis power switch to power down the system prevents remote power control using IPMI. Using the IPMI Remote System Power Down feature to power off the protection system does not perform an orderly shutdown.

For HA systems, a management connection to both nodes is required.

Steps

1. Run the following commands to verify that I/O on the system is stopped:
   - `cifs show active`
   - `nfs show active`
   - `system show stats view sysstat interval 2`
   - `system show perf`
2. For HA systems, run the `ha status` command to verify the health of the HA configuration.

   The following example is from a healthy system. If the system has a failed component, the HA System Status is degraded, and one or both nodes show offline for the HA State.

```
HA System Name: dd9900-ha3a.example.com
HA System Status: highly available
Node Name                  Node ID   Role      HA State
-------------------------- --------- --------- --------
dd9900-ha3a-p0.example.com   0       active    online
```

```
dd9900-ha3a-p1.example.com   1       standby   online
-------------------------- --------- --------- --------
```

3. Run the `alerts show current` command. For HA pairs, run the command on the active node first, and then the standby node.
4. For HA systems, run the `ha offline` command from the standby node if the system is in a highly available state with both nodes online. Skip this step if the HA status is degraded.
5. Run the `system poweroff` command. For HA pairs, run the command on the active node first, and then the standby node.
6. Remove the power cords from the power supplies on the controller or controllers.
7. Verify that the blue power LED on the controllers is off to confirm that the system is powered down.
8. When the controller has powered off, switch off any external expansion shelves.

Powering the DDOS system on

When powering a system on, follow the proper procedure to preserve the file system and configuration integrity.

About this task

Restore power to the protection system when the system downtime is complete.

Steps

1. Power on expansion shelves before powering on the controller. Wait approximately three minutes after all expansion shelves are turned on.

   NOTE: A controller includes the chassis and any internal storage. A DDOS system includes the controller and any external storage.

2. Plug in the power cord for the controller, and if there is a power button on the controller, press the power button, as shown in the *Installation Guide* for your system. For HA systems, power on the active node first, and then the standby node.

   NOTE: Some DDOS appliances do not have a traditional power button, and are designed to be "always on." These devices will power up as soon as AC power is applied.

   The system reboot time depends on the storage attached and might take approximately 30-40 minutes. Connect a console session to view the system boot sequence.
3. For HA systems, verify the health of the HA configuration.

   Run the command, `ha status`.
4. For HA systems, if one of the nodes displays as offline, run the `ha online` command on that node to restore the HA configuration.
   The `ha online` command triggers a system reboot.
5. Use a serial connection or an SSH session to verify that the system is fully booted and the operating system is running. The system is up when you can log into the system.
6. Run the `alerts show current` command. For HA pairs, run the command on the active node first, and then on the standby node.

Remote system power management with IPMI

Select DD systems support remote power management using the Intelligent Platform Management Interface (IPMI), and they support remote monitoring of the boot sequence using Serial over LAN (SOL).

IPMI power management takes place between an IPMI initiator and an IPMI remote host. The IPMI initiator is the host that controls power on the remote host. To support remote power management from an initiator, the remote host must be configured with an IPMI username and password. The initiator must provide this username and password when attempting to manage power on a remote host.

IPMI runs independently of DDOS and allows an IPMI user to manage system power as long as the remote system is connected to a power source and a network. An IP network connection is required between an initiator and a remote system. When

properly configured and connected, IPMI management eliminates the need to be physically present to power on or power off a remote system.

You can use both DD System Manager and the CLI to configure IPMI users on a remote system. After you configure IPMI on a remote system, you can use IPMI initiator features on another system to log in and manage power.

NOTE: If a system cannot support IPMI due to hardware or software limitations, DD System Manager displays a notification message when attempting to navigate to a configuration page.

SOL is used to view the boot sequence after a power cycle on a remote system. SOL enables text console data that is normally sent to a serial port or to a directly attached console to be sent over a LAN and displayed by a management host.

The DDOS CLI allows you to configure a remote system for SOL and view the remote console output. This feature is supported only in the CLI.

NOTE: IPMI power removal is provided for emergency situations during which attempts to shut down power using DDOS commands fail. IPMI power removal simply removes power to the system, it does not perform an orderly shutdown of the DDOS file system. The proper way to remove and reapply power is to use the DDOS `system reboot` command. The proper way to remove system power is to use the DDOS `system poweroff` command and wait for the command to properly shut down the file system.

Remove the power cords from the power supplies on the controller or controllers after powering the system down remotely.

IPMI and SOL limitations

IPMI amd SOL is supported on all systems supported by this release.

IPMI user support is as follows.:

- Maximum user IDs = 10.
- Two default users (NULL, root).
- Maximum user IDs available = 8.

Adding and deleting IPMI users with DD System Manager

Each system contains its own list of configured IPMI users, which is used to control access to local power management features. Another system operating as an IPMI initiator can manage remote system power only after providing a valid username and password.

About this task

This functionality is not supported on DD6900, DD9400, and DD9900 systems with DDOS 7.0 and later.

To give an IPMI user the authority to manage power on multiple remote systems, you must add that user to each of the remote systems.

NOTE: The IPMI user list for each remote system is separate from the DD System Manager lists for administrator access and local users. Administrators and local users do not inherit any authorization for IPMI power management.

Steps

1. Select **Maintenance** > **IPMI**.
2. To add a user, complete the following steps.
   a. Above the IPMI Users table, click **Add**.
   b. In the Add User dialog box, type the user name (16 or less characters) and password in the appropriate boxes (reenter the password in the **Verify Password** box).
   c. Click **Create**.

      The user entry appears in the **IPMI Users** table.
3. To delete a user, complete the following steps.
   a. In the IPMI Users list, select a user and click **Delete.**
   b. In the Delete User dialog box, click **OK** to verify user deletion.

Changing an IPMI user password

Change the IPMI user password to prevent use of the old password for power management.

About this task

This functionality is not supported on DD6900, DD9400, and DD9900 systems with DDOS 7.0 and later.

Steps

1. Select **Maintenance** > **IPMI**.
2. In the IPMI Users table, select a user, and click **Change Password**.
3. In the Change Password dialog box, type the password in the appropriate text box and reenter the password in the **Verify Password** box.
4. Click **Update**.

Configuring an IPMI port

When you configure an IPMI port for a system, you select the port from a network ports list and specify the IP configuration parameters for that port. The selection of IPMI ports displayed is determined by the protection system model.

About this task

Some systems support one or more dedicated ports, which can be used only for IPMI traffic. Other systems support ports that can be used for both IPMI traffic and all IP traffic supported by the physical interfaces in the **Hardware** > **Ethernet** > **Interfaces** view. Shared ports are not provided on systems that provide dedicated IPMI ports.

The port names in the IPMI Network Ports list use the prefix bmc, which represents baseboard management controller. To determine if a port is a dedicated port or shared port, compare the rest of the port name with the ports in the network interface list. If the rest of the IPMI port name matches an interface in the network interface list, the port is a shared port. If the rest of the IPMI port name is different from the names in the network interface list, the port is a dedicated IPMI port.

When IPMI and nonIPMI IP traffic share an Ethernet port, if possible, do not use the link aggregation feature on the shared interface because link state changes can interfere with IPMI connectivity.

Steps

1. Select **Maintenance** > **IPMI**.

   The IPMI Configuration area shows the IPMI configuration for the managed system. The Network Ports table lists the ports on which IPMI can be enabled and configured. The IPMI Users table lists the IPMI users who can access the managed system.

2. In the **Network Ports** table, select a port to configure.

   NOTE: If the IPMI port also supports IP traffic (for administrator access or backup traffic), the interface port must be enabled before you configure IPMI.

3. Above the **Network Ports** table, click **Configure**.

   The Configure Port dialog box appears.

4. Choose how network address information is assigned.
   - To collect the IP address, netmask, and gateway configuration from a DHCP server, select **Dynamic (DHCP)**.
   - To manually define the network configuration, select **Static (Manual)** and enter the IP address, netmask, and gateway address.
   - Click **Apply**.
5. Enable a disabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Enable.**
6. Disable a disabled IPMI network port by selecting the network port in the **Network Ports** table, and clicking **Disable.**

Preparing for remote power management and console monitoring with the CLI

Remote console monitoring uses the Serial Over Lan (SOL) feature to enable viewing of text-based console output without a serial server. You must use the CLI to set up a system for remote power management and console monitoring.

About this task

Remote console monitoring is typically used in combination with the `ipmi remote power cycle` command to view the remote system's boot sequence. This procedure should be used on every system for which you might want to remotely view the console during the boot sequence.

Steps

1. Connect the console to the system directly or remotely.
   - Use the following connectors for a direct connection.
     - DIN-type connectors for a PS/2 keyboard
     - USB-A receptacle port for a USB keyboard
     - DB15 female connector for a VGA monitor
   - For a serial connection, use a standard DB9 male or micro-DB9 female connector. A null modem cable with male micro-DB9 and standard female DB9 connectors is included for a typical laptop connection.
   - For a remote IPMI/SOL connection, use the appropriate RJ45 receptacle as follows.
     - For other systems, use the maintenance or service port. For port locations, refer to the system documentation, such as a hardware overview or installation and setup guide.
2. To support remote console monitoring, use the default BIOS settings.
3. To display the IPMI port name, enter `ipmi show config`.
4. To enable IPMI, enter `ipmi enable {`*port* `| all}`.
5. To configure the IPMI port, enter `ipmi config port { dhcp | ipaddress` *ipaddr* `netmask` *mask* `gateway` *ipaddr* `}`.

   NOTE: If the IPMI port also supports IP traffic (for administrator access or backup traffic), the interface port must be enabled with the `net enable` command before you configure IPMI.
6. If this is the first time using IPMI, run `ipmi user reset` to clear IPMI users that may be out of synch between two ports, and to disable default users.
7. To add a new IPMI user, enter `ipmi user add` *user*.
8. To set up SOL, do the following:
   a. Enter `system option set console lan`.
   b. When prompted, enter **`y`** to reboot the system.

Managing power with DD System Manager

After IPMI is properly set up on a remote system, you can use DD System Manager as an IPMI initiator to log into the remote system, view the power status, and change the power status.

Steps

1. Select **Maintenance** > **IPMI**.
2. Click **Login to Remote System**.

   The IPMI Power Management dialog box appears.
3. Enter the remote system IPMI IP address or hostname and the IPMI username and password, then click **Connect**.
4. View the IPMI status.

   The IPMI Power Management dialog box appears and shows the target system identification and the current power status. The Status area always shows the current status.

NOTE: The Refresh icon (the blue arrows) next to the status can be used to refresh the configuration status (for example, if the IPMI IP address or user configuration were changed within the last 15 minutes using the CLI commands).

5. To change the IPMI power status, click the appropriate button.
   - **Power Up**—Appears when the remote system is powered off. Click this button to power up the remote system.
   - **Power Down**—Appears when the remote system is powered on. Click this button to power down the remote system.
   - **Power Cycle**—Appears when the remote system is powered on. Click this button to power cycle the remote system.
   - **Manage Another System**—Click this button to log into another remote system for IPMI power management.
   - **Done**—Click to close the IPMI Power Management dialog box.

   NOTE: The IPMI Power Down feature does not perform an orderly shutdown of the DDOS. This option can be used if the DDOS hangs and cannot be used to gracefully shutdown a system.

Managing power with the CLI

You can manage power on a remote system and start remote console monitoring using the CLI.

About this task

NOTE: The remote system must be properly set up before you can manage power or monitor the system.

Steps

1. Establish a CLI session on the system from which you want to monitor a remote system.
2. To manage power on the remote system, enter `ipmi remote power {on | off | cycle | status} ipmi-target <ipaddr | hostname> user user`.
3. To begin remote console monitoring, enter `ipmi remote console ipmi-target <ipaddr | hostname> user user`.

   NOTE: The user name is an IPMI user name defined for IPMI on the remote system. DDOS user names are not automatically supported by IPMI.

4. To disconnect from a remote console monitoring session and return to the command line, enter the at symbol (@).
5. To terminate remote console monitoring, enter the tilde symbol (~).

Use iDRAC to power the system on and off remotely

About this task

This task applies to DD3300, DD6400, DD6900, DD9400, and DD9900 systems only.

Steps

1. In a web browser, type the iDRAC IP address specified during iDRAC configuration.
2. Login with the user name root. The default password is the system serial number on the PSNT.
3. Select **Dashboard**.
4. Click **Graceful Shutdown** to initiate the same behavior as pressing the power button, or select the drop-down arrow to select on the of the following options:
   - **Power Off System**
   - **Reset System (warm boot)**
   - **Power Cycle System (cold boot)**

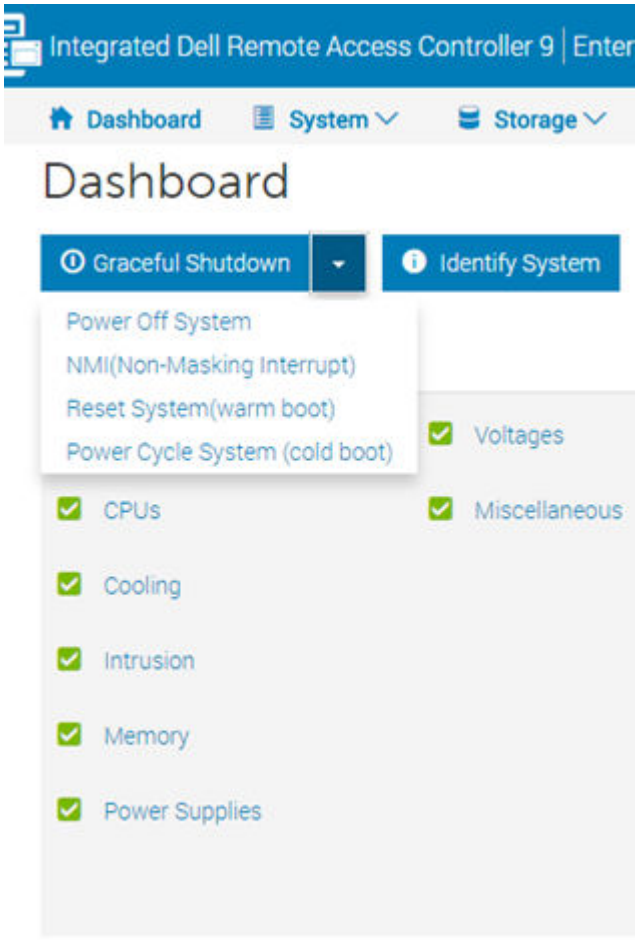

Figure 2. iDRAC power options

5. Remove the power cords from the power supplies on the controller or controllers after powering the system down from iDRAC.

6

Setting up Support

This chapter presents the following topics:

Topics:

Managing system support

The **Maintenance** > **Support** view displays tabs for support configuration, autosupport reporting, support bundle management, and core dump management.

Managing HA system autosupport and support bundles

The standby node mirrors the configuration on the active node, but the ASUP and support bundle for each node is different.

Contents of the Autosupport and support bundles for each node are:

- Active node—Local node, file system, replication, protocol, and full HA information
- Standby node—Local node information and some HA configuration and status information

Autosupport and support bundles from both the nodes are needed to debug issues related to HA system status (filesystem, replication, protocols, and HA configuration).

Managing the support configuration

Use the **Configuration** tab to setup Secure Connect Gateway (SCG), CloudIQ, alert reporting, and email subscription lists for autosupport reporting and alert summaries.

For HA pairs, configure support from the active node. The configuration controls are disabled on the standby node.

Configure the support channel

The support channel is the method by which system Autosupport, alert summaries, and real time alerts are shared with Dell Support. By default, this is set to email information to Dell Support. Configuring Secure Connect Gateway (SCG) instead of email delivery enables remote assistance from Dell Support, and CloudIQ reporting.

Prerequisites

Steps

1. Select **Maintenance** > **Support** > **Configuration**.
2. In the **Support Channel** pane, click **Configure SCG**.
3. In the **Gateway Hostname** field, specify the hostname of the SCG gateway.
4. In the **DD System Local IP** list box, select the IP address of the DD system.

5. In the **Username** field, specify the service link username.
6. In the **Password** field, specify the service link password.
7. Click **Configure**.

Next steps

Click **Edit** to modify the support channel configuration if changes are required later.

Enabling and disabling CloudIQ data sharing

CloudIQ is a no cost cloud-native application that leverages Machine Learning to proactively monitor and measure the overall health of storage systems through intelligent, comprehensive, and predictive analytics. DDOS versions 7.4 and higher are CloudIQ-ready. The Cloud IQ release updates provide more information about CloudIQ support for Data Domain and PowerProtect DD monitoring.

Prerequisites

CloudIQ data sharing requires the selection of Secure Connect Gateway as the support channel.

About this task

The **Learn More** link in the CloudIQ area provides additional information about CloudIQ.

Steps

1. Select **Maintenance** > **Support** > **Configuration**.
2. In the **CloudIQ** pane, click **Enable**.
3. Click **Enable**.
4. Log into CloudIQ at https://cloudiq.dellemc.com/ with your Dell credentials.
5. Review the welcome screens, and click **Continue** to allow CloudIQ to access information about the user and systems associated with the specified Dell account.
6. Click **Continue**.
7. If CloudIQ does not detect that it is receiving data from the PowerProtect DD system, select **Admin** > **Connectivity** to view the detected CloudIQ-compatible systems and verify that the DD system is configured to use Secure Remote Services and to send data to CloudIQ.

Next steps

Click **Disable** to disable CloudIQ reporting if required at a later time.

Configure telemetry

Specify the settings, details, and schedules for sharing system alert summaries and autosupport reports, and enable or disable sharing that data with Dell Support.

Steps

1. Select **Maintenance** > **Support** > **Configuration**.
2. In the **Telemetry** pane, click **Configure**.
3. Specify the alert summary details.
   a. Select the **Enable** checkbox to enable system alert summaries.
   b. Select a **Daily** or **Weekly** schedule.
   c. Specify a delivery time.
   d. Select **Text** or **HTML** email formatting.
   e. Click **Next**.
4. Specify the autosupport details.
   a. Select the **Enable** checkbox to enable system autosupport reports.
   b. Select a **Daily** or **Weekly** schedule.

c. Specify a delivery time.
d. Click **Next**.

5. In the **Subject Tag** field, optionally specify additional text to append to the subject of the alert summary and autosupport emails.

   NOTE: This field has a maximum length of 64 characters.

   a. Click **Next**.
6. Select the **Enable** checkboxes as required to specify what information is shared with Dell Support.
   a. Click **Next**.
7. Review the summary and click **Finish**.

Next steps

Click **Configure** to modify the telemetry configuration if changes are required later.

Configure autosupport subscribers

Create an email distribution list for system autosupport reports. The system administrator email address is automatically added to the list.

Steps

1. Select **Maintenance** > **Support** > **Configuration**.
2. In the **Autosupport Subscribers** pane, click **Add**.
3. Specify the email address to add.

Next steps

Click **Remove** next to an email address to remove it from the list at a later time, or click.**Add** to add a new email address to the list.

Configure alert summary subscribers

Create an email distribution list for system alert summaries. The system administrator email address is automatically added to the list.

Steps

1. Select **Maintenance** > **Support** > **Configuration**.
2. In the **Alert Summary Subscribers** pane, click **Add**.
3. Specify the email address to add.

Next steps

Click **Remove** next to an email address to remove it from the list at a later time, or click.**Add** to add a new email address to the list.

Managing Autosupport reports

The Autosupport feature generates a report that is called an Auto Support log (ASUP). The ASUP shows system identification information, consolidated output from several system commands, and entries from various log files. Extensive and detailed internal statistics appear at the end of the report. This report is designed to aid Support in debugging system problems.

An ASUP is generated as scheduled, which is usually once per day. Additionally, every time the file system starts, the system generates a new ASUP.

Other reports are triggered by system events such as alerts, and are more limited in scope. They contain basic system information, and information about the event that triggered the report.

Reviewing generated autosupport reports

Review autosupport reports to view system statistics and configuration information captured in the past. The system stores a maximum of 14 autosupport reports.

Steps

Select **Maintenance** > **Support** > **Autosupport Reports**.

The Autosupport Reports page shows the autosupport report file name and file size, and the date the report was generated. Reports are automatically named. The most current report is autosupport, the previous day is autosupport.1, and the number increments as the reports move back in time.

Managing support bundles

A support bundle is a file that contains system configuration and operation information. It is a good practice to generate a support bundle before a software upgrade or a system topology change (such as a controller upgrade).

Dell Support often requests a support bundle when providing assistance.

The KB articles *Data Domain: How to collect/upload a support bundle (SUB) from a Data Domain Restorer (DDR)* and *Data Domain: Gathering Autosupports*, available from the Online Support website, provide additional information about gathering and working with support bundles.

Generating a support bundle

When troubleshooting problems, Dell Support may ask for a support bundle, which is a tar-g-zipped selection of log files with a README file that includes identifying autosupport headers.

Steps

1. Select **Maintenance** > **Support** > **Support Bundles**.
2. Click **Generate Support Bundle**.

   NOTE: The system supports a maximum of three support bundles. If you attempt to generate a fourth support bundle, the system automatically deletes the oldest support bundle. You can also delete support bundles using the CLI command `support bundle delete`.

   Also, if you generate a support bundle on a upgraded system that contains a support bundle named using the old format, *`hostname`*`-support-bundle-`*`timestamp`*`.tar.gz`, that file is renamed to use the newer name format.

3. In the **Bundle Type** list box, select **Mini Bundle** or **Full Bundle**.
4. In the **Duration** list box, select a duration to include in the support bundle.
   - **All**
   - **Last *N* days**
   - **Custom date range**
5. If required, specify a number of days (**Last *N* days**) or start and end dates (**Custom date range**).
6. Click **Generate**.
7. Email the file to customer support at support@emc.com.

   NOTE: If the bundle is too large to be emailed, use the online support site to upload the bundle. (Go to https://www.dell.com/support.)

Managing a core dump

A core is a file that contains details about the specific problem encountered when the protection system suffers a crash due to a core dump. DDOS keeps a record of these files to assist with troubleshooting.

Navigate to **Maintenance** > **Support** > **Cores**.

If a core file is too big, DDOS provides the ability to split it into smaller chunks. Split files are automatically deleted after 48 hours.

Coredump analysis

DDOS provides the ability to automatically triage coredump files and provide a summary report of the information contained in the coredump. The reports are stored in a protected directory under `/ddvar`, and are accessible by any administrative access method that requires sysadmin authentication such as FTP or SCP. Use the `support coredump analysis enable` and `support coredump analysis disable` commands to turn this functionality on or off.

Additionally, the `support notification enable` command provides an option to include the coredump analysis in the telemetry the system sends to Dell.

Splitting a coredump file

When DDOS crashes due to a coredump, a core file describing the problem is created in the `/ddvar/core` directory. This file may be large, and difficult to copy off the protection system. If the core file cannot be copied off the system because it is too large, DDOS provides the ability to split the core file into smaller chunks.

Steps

1. Select **Maintenance** > **Support** > **Cores**.
2. Select a core file from the table.
3. Click **Split**.
4. In the **Size** field, specify the size of the chunks to create and select **MiB** or **GiB** from the list box.

   NOTE: A single core file can be broken down into a maximum of 20 chunks. The command will fail with an error if the specified size would result in more than 20 chunks.

5. Click **OK**.

Results

DDOS splits the selected coredump file into chunks of the specified size, and places them in the `/ddvar/core` directory. Split files are automatically deleted after 48 hours.

CLI equivalent

Steps

1. Run the **`support coredump split <filename> <n> {MiB|GiB}`** command, where:
   - *<filename>* is the name of the core file in the `/ddvar/core` directory
   - *<n>* is the size of the smaller chunks to create

     NOTE: A single core file can be broken down into a maximum of 20 chunks. The command will fail with an error if the specified size would result in more than 20 chunks.

   For example, splitting a 42.1 MB core file named `cpmdb.core.19297.1517443767` into 10 MB chunks would result in five chunks.

```
# support coredump split cpmdb.core.19297.1517443767 10 MiB
cpmdb.core.19297.1517443767 will be split into 5 chunks.
Splitting...

The md5 and split chunks of cpmdb.core.19297.1517443767:
File                                Size       Time Created
---------------------------------   --------   ------------------------
cpmdb.core.19297.1517443767_5_01    10.0 MiB   Mon Feb  5 11:50:57 2018
cpmdb.core.19297.1517443767_5_02    10.0 MiB   Mon Feb  5 11:50:57 2018
cpmdb.core.19297.1517443767_5_03    10.0 MiB   Mon Feb  5 11:50:57 2018
cpmdb.core.19297.1517443767_5_04    10.0 MiB   Mon Feb  5 11:50:57 2018
cpmdb.core.19297.1517443767_5_05    2.1 MiB    Mon Feb  5 11:50:57 2018
cpmdb.core.19297.1517443767.md5     0 MiB      Mon Feb  5 11:50:58 2018
```

```
------------------------------   --------   ------------------------
Download the files as soon as possible. Otherwise they will be automatically delete in
48 hours.
```

2. Run the **support coredump save *<file-list>*** command to save specified coredump files to a USB drive. Split files are automatically deleted after 48 hours.

Managing log files

The DD system maintains a set of log files, which can be bundled and sent to Support to assist in troubleshooting any system issues that may arise. Log files cannot be modified or deleted by any user with DD System Manager, but they can be copied from the log directory and managed off of the system.

NOTE: Log messages on an HA system are preserved on the node where the log file originated.

Log files are rotated weekly. Every Sunday at 0:45 a.m., the system automatically opens new log files for the existing logs and renames the previous files with appended numbers. For example, after the first week of operation, the previous week `messages` file is renamed `messages.1`, and new messages are stored in a new messages file. Each numbered file is rolled to the next number each week. For example, after the second week, the file `messages.1` is rolled to `messages.2`. If a `messages.2` file already existed, it rolls to `messages.3`. At the end of the retention period, the expired log is deleted. For example, an existing `messages.9` file is deleted when `messages.8` rolls to `messages.9`.

The `audit.log` does not rotate on a weekly basis. Instead, it rotates when the file reaches 70 MB in size.

Except as noted in this topic, the log files are stored in `/ddvar/log`.

NOTE: Files in the `/ddvar` directory can be deleted using Linux commands if the Linux user is assigned *write* permission for that directory.

The set of log files on each system is determined by the features configured on the system and the events that occur.

By default, the system uses port 514 for log files. To select a different port, run the **log server-port set *<port-number>*** command, and specify any unused port on the system up to a maximum of 65535.

Viewing log files in DD System Manager

Use the Logs tab to view and open the system log files in DD System Manager.

Steps

1. Select **Maintenance** > **Logs**.

   The Logs list displays log file names and the size and generation date for each log file.

2. Click a log file name to view its contents. You may be prompted to select an application, such as Notepad.exe, to open the file.

Displaying a log file in the CLI

Use the `log view` command to view a log file in the CLI.

Steps

1. To view a log file in the CLI, use the `log view` command.

   With no argument, the command displays the current messages file.

2. When viewing the log, use the up and down arrows to scroll through the file; use the q key to quit; and enter a slash character (`/`) and a pattern to search through the file.

Example

The display of the messages file is similar to the following. The last message in the example is an hourly system status message that the protection system generates automatically. The message reports system uptime, the amount of data stored, NFS operations, and the amount of disk space used for data storage (%). The hourly messages go to the system log and to the serial console if one is attached.

```
# log view
Jun 27 12:11:33 localhost rpc.mountd: authenticated unmount request from perfsun-
g.emc.com:668 for /ddr/col1/segfs (/ddr/col1/segfs)

Jun 27 12:28:54 localhost sshd(pam_unix)[998]: session opened for user jsmith10 by (uid=0)

Jun 27 13:00:00 localhost logger: at 1:00pm up 3 days, 3:42, 52324 NFS ops, 84763 GiB data
col. (1%)
```

NOTE: GiB = Gibibytes = the binary equivalent of Gigabytes.

Learning more about log messages

Look up error messages in the Error Message Catalog for your DDOS version.

About this task

In the log file is text similar to the following.

```
Jan 31 10:28:11 syrah19 bootbin: NOTICE: MSG-SMTOOL-00006: No replication throttle
schedules found: setting throttle to unlimited.
```

The components of the message are as follows.

```
DateTime Host Process [PID]: Severity: MSG-Module-MessageID: Message
```

Severity levels, in descending order, are: Emergency, Alert, Critical, Error, Warning, Notice, Info, Debug.

Steps

1. Go to the Online Support website at https://www.dell.com/support, enter *Error Message Catalog* in the search box, and click the search button.
2. In the results list, locate the catalog for your system and click on the link.
3. User your browser search tool to search for a unique text string in the message.
   The error message description looks similar to the following display.

   ```
   ID: MSG-SMTOOL-00006 - Severity: NOTICE - Audience: customer

   Message: No replication throttle schedules found: setting throttle to unlimited.

   Description: The restorer cannot find a replication throttle schedule. Replication is
   running with throttle set to unlimited.

   Action: To set a replication throttle schedule, run the replication throttle add command.
   ```

4. To resolve an issue, do the recommended action.
   Based on the example message description, one could run the `replication throttle add` command to set the throttle.

Saving a copy of log files

Save log file copies to another device when you want to archive those files.

About this task

Use NFS, CIFS mount, or FTP to copy the files to another machine. If using CIFS or NFS, mount /ddvar to your desktop and copy the files from the mount point. The following procedure describes how to use FTP to move files to another machine.

Steps

1. On the protection system, use the `adminaccess show ftp` command to see whether FTP service is enabled. If the service is disabled, use the command `adminaccess enable ftp`.
2. Use the `adminaccess show ftp` command to see that the FTP access list includes the IP address of your remote machine. If the address is not in the list, use the command `adminaccess add ftp` *ipaddr*.
3. On the remote machine, open a web browser.
4. In the **Address** box at the top of the web browser, use FTP to access the protection system as shown in the following example.

   `ftp://Data Domain system_name.yourcompany.com/`

   NOTE: Some web browsers do not automatically ask for a login if a machine does not accept anonymous logins. In that case, add a user name and password to the FTP line. For example: **`ftp://sysadmin:your-pw@Data Domain system_name.yourcompany.com/`**

5. At the login pop-up, log into the protection system as user **`sysadmin`**.
6. On the protection system, you are in the directory just above the log directory. Open the log directory to list the messages files.
7. Copy the file that you want to save. Right-click the file icon and select **Copy To Folder** from the menu. Choose a location for the file copy.
8. If you want the FTP service disabled on the protection system, after completing the file copy, use SSH to log into the protection system as sysadmin and invoke the command `adminaccess disable ftp`.

Log message transmission to remote systems

Some log messages can be sent from the protection system to other systems. DDOS uses syslog to publish log messages to remote systems.

A protection system exports the following facility.priority selectors for log files. For information on managing the selectors and receiving messages on a third-party system, see your vendor-supplied documentation for the receiving system.

- *.notice—Sends all messages at the notice priority and higher.
- *.alert—Sends all messages at the alert priority and higher (alerts are included in *.notice).
- kern.*—Sends all kernel messages (kern.info log files).

The `log host` commands manage the process of sending log messages to another system.

Viewing the log file transmission configuration

Use the `log host show` CLI command to view whether log file transmission is enabled and which hosts receive log files.

Steps

To display the configuration, enter the `log host show` command.

Example

```
# log host show
Remote logging is enabled.
Remote logging hosts
     log-server
```

Enabling and disabling log message transmission

You must use CLI commands to enable or disable log message transmission.

Steps

1. To enable sending log messages to other systems, use the `log host enable` command.
2. To disable sending log messages to other systems, use the `log host disable` command.

Adding or removing a receiver host

You must use CLI commands to add or remove a receiver host.

Steps

1. To add a system to the list that receives protection system log messages, use the `log host add` command.
2. To remove a system from the list that receives system log messages, use the command: `log host del`.

Example

The following command adds the system named *log-server* to the hosts that receive log messages.

```
log host add log-server
```

The following command removes the system named *log-server* from the hosts that receive log messages.

```
log host del log-server
```

The following command disables the sending of logs and clears the list of destination hostnames..

```
log host reset
```

Configuring secure-syslog for encrypted log forwarding

DDOS provides the ability to encrypt log forwarding to a remote host. Use the CLI to configure this functionality.

About this task

Secure-syslog supports anonymous mode, which only uses server side certificate authentication. This requires importing:

- The syslog server CA certificate on the DD system
- The host certificate and host key on the syslog server system

The default secure-syslog server port is 10514. DDOS supports multiple syslog servers. When multiple syslog servers are configured, they use the same secure-syslog server port configured on the DD system.

The system raises an alert when log forwarding through secure-syslog fails. The alert conditions are:

- Connection to the secure-syslog server fails
- The secure-syslog server CA certificate appears to be invalid

Complete the following steps to configure secure-syslog.

Steps

1. Run the **`log secure-syslog host add <host>`** command to add the secure-syslog host to the system.

```
# log secure-syslog host add 10.198.177.6
Secure remote host logging is not currently enabled. Enable with 'log secure-syslog host
enable'.
Host "10.198.177.6" added.
```

2. Run the **`adminaccess certificate import ca application secure-syslog`** command to import the CA certificate for the secure-syslog server.

```
# adminaccess certificate import ca application secure-syslog
Enter the certificate and then press Control-D, or press Control-C to cancel.

MIIDgTCCAmmgAwIBAgIJAIsFi6huU/QSMA0GCSqGSIb3DQEBCwUAMFcxCzAJBgNV
BAYTAlVTMQswCQYDVQQIDAJZWTELMAkGA1UEBwwCSlMxHDAaBgNVBAoME0RlZmF1
bHQgQ29tcGFueSBMdGQxEDAOBgNVBAMMBzEuMi4zLjQwHhcNMjMwMzI4MDc0MzAz
WhcNMjYwMzI3MDc0MzAzWjBXMQswCQYDVQQGEwJVUzELMAkGA1UECAwCWVkxCzAJ
BgNVBAcMAkpTMRwwGgYDVQQKDBNEZWZhdWx0IENvbXBhbnkgTHRkMRAwDgYDVQQD
DAcxLjIuMy40MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEArQIWAlhv
KqSY/iaXS5O6vxJ9HHSgts7OWI/uhjO9yA/may2yHBvjxmLheg1ixseOjpKxvLfY
Of9ufLGKWpbVIJGoXkG6x+zde1hwbctK4EhN0XTJ/xoUVVu/F2DqeeM1B6bt+26Q
GR2xx3kJuFMBxtDvcrql/yXPH2BNPHyJJ6CIa1hwbx5iwxJJNUkLe/pjKBhRNyS0
T4trEwGgsNOVSyYGCkAo3BWPyijBagatPQFs36SrOVc3AcFu3ie9q67NEJDxOwfk
```

```
iFHvT8zYVRkDNYYmN7wt76TGK4G8HuldyZ19z+0fa6m6pMKbuOht19PtiP3MBXWh
e+jZYuzrmPVRHwIDAQABo1AwTjAdBgNVHQ4EFgQUXLSU4KHIiNlnXAKJdCexeA9X
ROwwHwYDVR0jBBgwFoAUXLSU4KHIiNlnXAKJdCexeA9XROwwDAYDVR0TBAUwAwEB
/zANBgkqhkiG9w0BAQsFAAOCAQEAUt0kgFbSfegkskrVDv4DwKKWlIkxgBJEVsvH
y+T16KszhedvUOUIM2quv6J1E1BqmrUlQSYb8RbJqOO6vWpruxVd4RYBSRIJzQT0
p3fGV3M90oi/bhmSt7v/Q7DpzzJgxDVSuKNXMf4WgPY212pubmUMfJFkDkK0t/pG
5OnLL9ChsAvZSX5mHDr7wbojO+GJJNAeLvLSBVtnNyB1e1xj0dpheIYyVP329sPN
C79uP+HdXma0ujOQgqpnpwAYY0faB6tcb/mkn/SyL30Fx01HaXRwdF6CivoakOgw
Hkrf88XDMPXBK4kstEqGoO0RRFPL0tAQN4hu+hQpRmr03nzhyQ==

The SHA1 fingerprint for the imported CA certificate is:
AD:61:28:84:71:EB:5B:7F:E7:9A:EC:3B:16:25:9B:99:28:9E:33:58

        Do you want to import this certificate? (yes|no) [yes]: yes
CA certificate imported for application(s) : "secure-syslog".
```

3. Run the **adminaccess certificate show imported-ca application secure-syslog** to verify the certificate was imported.

```
# adminaccess certificate show imported-ca application secure-syslog

Subject   Type          Application     Valid From                Valid Until
     Fingerprint
-------   -----------   -------------   ------------------------
------------------------   ------------------------------------------------------------
1.2.3.4   imported-ca   secure-syslog   Tue Mar 28 00:43:03 2023   Fri Mar 27 00:43:03
2026   AD:61:28:84:71:EB:5B:7F:E7:9A:EC:3B:16:25:9B:99:28:9E:33:58
-------   -----------   -------------   ------------------------
------------------------   ------------------------------------------------------------
Certificate signing request (CSR) exists at /ddvar/certificates/
CertificateSigningRequest.csr
```

4. Run the **log secure-syslog host enable** command to enable secure-syslog log forwarding.

```
# log secure-syslog host enable
Secure-syslog remote host logging enabled.
```

5. Run the **log secure-syslog host show** command to verify secure-syslog log forwarding is enabled.

```
# log secure-syslog host show
Secure-syslog remote logging is enabled.
Remote logging hosts
    10.198.177.6
```

6. Run the **log secure-syslog server-port show** command to check the port for secure-syslog log forwarding.

```
# log secure-syslog server-port show
Server-port     10514
```

7. If necessary, run the **log secure-syslog server-port set *port-number*** command to change the port number.

Next steps

The syslog server requires a CA certificate, host certificate, and host key. The following example shows a sample server-side configuration for secure-syslog:

```
global(
DefaultNetstreamDriver="gtls"
DefaultNetstreamDriverCAFile="/etc/rsyslog.d/cert/cacert.pem"
DefaultNetstreamDriverCertFile="/etc/rsyslog.d/cert/ser_cert.pem"
DefaultNetstreamDriverKeyFile="/etc/rsyslog.d/cert/serkey.pem"
)



$ModLoad imtcp # TCP listener

$InputTCPServerStreamDriverMode 1 # run driver in TLS-only mode
$InputTCPServerStreamDriverAuthMode anon
```

```
$InputTCPServerRun 10514 # start up listener at port 10514
```

7

Managing Alerts

This chapter presents the following topics:

Topics:

Health Alerts panel

Alerts are messages from system services and subsystems that report system events. The Health > Alerts panel displays tabs that allow you to view current and non-current alerts, the configured alert notification groups, and the configuration for those who want to receive daily alert summary reports.

Alerts are also sent as SNMP traps. See the *MIB Quick Reference Guide* or the SNMP MIB for the full list of traps.

Related concepts

Managing alert notifications

Viewing and clearing current alerts

The Current Alerts tab displays a list of all the current alerts and can display detailed information for a selected alert. An alert is automatically removed from the Current Alerts list when the underlying situation is corrected or when manually cleared.

Steps

1. To view all of the current alerts, select **Health** > **Alerts** > **Current Alerts**.
2. To limit the number of entries in the current alert list, do the following.
   a. In the Filter By area, select a **Severity** and **Class** to expose only alerts that pertain to those choices.
   b. Click **Update**.

      All alerts not matching the Severity and Class are removed from the list.
3. To display additional information for a specific alert in the **Details** area, click the alert in the list.
4. To clear an alert, select the alert checkbox in the list and click **Clear.**

   A cleared alert no longer appears in the current alerts list, but it can be found in the alerts history list.
5. To remove filtering and return to the full listing of current alerts, click **Reset**.

Related concepts

Managing alert notifications

Viewing the alerts history

The Alerts History tab displays a list of all the cleared alerts and can display detailed information for a selected alert.

Steps

1. To view all of the alerts history, select **Health** > **Alerts** > **Alerts History**.
2. To limit the number of entries in the current alert list, do the following.
   a. In the Filter By area, select a **Severity** and **Class** to expose only alerts that pertain to those choices.
   b. Click **Update**.

      All alerts not matching the Severity and Class are removed from the list.
3. To display additional information for a specific alert in the **Details** area, click the alert in the list.
4. To remove filtering and return to the full listing of cleared alerts, click **Reset**.

Related concepts

Managing alert notifications

Managing alert notifications

The alert feature generates event and summary reports that can be distributed to configurable email lists and to Dell.

Event reports are sent immediately and provide detailed information on a system event. The distribution lists for event alerts are called *notification groups*. You can configure a notification group to include one or more email addresses, and you can configure the types and severity level of the event reports sent to those addresses. For example, you might configure one notification group for individuals who need to know about critical events and another group for those who monitor less critical events. Another option is to configure groups for different technologies. For example, you might configure one notification group to receive email messages about all network events and another group to receive messages about storage issues.

Summary reports are sent daily and provide a summary of the events that occurred during the last 24 hours. Summary reports do not include all the information that is provided in event reports. The default generation time for the daily report is 08.00 a.m, and it can be changed. Summary reports are sent using a dedicated email list that is separate from the event notification groups.

You can enable or disable alert distribution to Dell. When sending reports to Dell, you have the option to select the legacy unsecure method or Secure Connect Gateway for secure transmissions.

HA system alert notification management

The alert feature on an HA system generates event and summary report like a non-HA system but how the HA system manages these alerts is different due to the two node system set-up.

Initial alert configuration is completed on the active node and mirrored to the stand-by (i.e., same configuration on both nodes). Local and AM-Alerts are emailed according to the notification settings and include information indicating they are from an HA system and from which node, the active or standby, that generated the alerts.

If there are active alerts on the file system, replication, or protocols when a failover occurs, these active alerts continue to show on the new active node after failover if the alert conditions have not cleared up.

Historical alerts on the filesystem, replication, and protocols stay with the node where they originated rather than failing over together with the filesystem on a failover. This means the CLIs on the active node will not present a complete/continuous view of historical alerts for filesystem, replication, and protocols

During a failover, local historical alerts stay with the node from which they were generated; however, the historical alerts for the filesystem, replication, and protocols (generally called "logical alerts") fail over together with the filesystem.

NOTE: The **Health** > **High Availability** panel displays only alerts that are HA-related. Those alerts can be filtered by major HA component, such as HA Manager, Node, Interconnect, Storage, and SAS connection.

Viewing the notification group list

A notification group defines a set of alert types (classes) and a group of email addresses (for subscribers). Whenever the system generates an alert type selected in a notification list, that alert is sent to the list subscribers.

Steps

1. Select **Health** > **Alerts** > **Notification**.
   CLI equivalent

   ```
   # alerts notify-list show
   ```

2. To limit (filter) the entries in the Group Name list, type a group name in the Group Name box or a subscriber email in the Alert Email box, and click **Update**.

   NOTE: Click **Reset** to display all configured groups.

3. To display detailed information for a group, select the group in the Group Name list.

Creating a notification group

Use the Notification tab to add notification groups and select the severity level for each group.

Steps

1. Select **Health** > **Alerts** > **Notification**.
2. Click **Add**.

   The Add Group dialog box appears.

3. Type the group name in the **Group Name** box.
4. Select the checkbox of one or more alert classes of which to be notified.
5. To change the default severity level (Warning) for a class, select another level in the associated list box.

   The severity levels are listed in ascending severity level. *Emergency* is the highest severity level.
6. Click **OK**.
   CLI equivalent

   ```
   # alerts notify-list create eng_grp class hardwareFailure
   ```

Managing the subscriber list for a group

Use the Notification tab to add, modify, or delete email addresses from a notification group subscriber list.

Steps

1. Select **Health** > **Alerts** > **Notification**.
2. Select the checkbox of a group in the Notifications group list, and do one of the following.
   - Click **Modify** and select **Subscribers.**
   - Click **Configure** in the Subscribers list.
3. To add a subscriber to the group, do the following.
   a. Click the **+** icon.

      The Email Address dialog box appears.

   b. Enter the email address of a subscriber.
   c. Click **OK**.
      CLI equivalent

```
# alerts notify-list add eng_lab emails mlee@urcompany.com,bob@urcompany.com
```

4. To modify an email address, do the following.
   a. Click the checkbox of the email address in the **Subscriber Email** list.
   b. Click the pencil icon.
   c. Edit the email address in the Email Address dialog box.
   d. Click **OK**.
5. To delete an email address, click the checkbox of the email address in the **Subscriber Email** list and click the **X** icon.

   CLI equivalent

```
# alerts notify-list del eng_lab emails bob@urcompany.com
```

6. Click **Finish** or **OK**.

Modifying a notification group

Use the Notification table to modify the attribute classes in an existing group.

Steps

1. Select **Health** > **Alerts** > **Notification**.
2. Select the checkbox of the group to modify in the group list.
3. To modify the class attributes for a group, do the following.
   a. Click **Configure** in the Class Attributes area.
   The Edit Group dialog box appears.
   b. Select (or clear) the checkbox of one or more class attributes.
   c. To change the severity level for a class attribute, select a level from the corresponding list box.
   d. Click **OK**.

   CLI equivalent

```
# alerts notify-list add eng_lab class cloud severity warning
# alerts notify-list del eng_lab class cloud severity notice
```

4. To modify the subscriber list for a group, do the following.
   a. Click **Configure** in the Subscribers area.
   The Edit Subscribers dialog box appears.
   b. To delete subscribers from the group list, select the checkboxes of subscribers to delete and click the Delete icon (X).
   c. To add a subscriber, click the Add icon (+), type a subscriber email address, and click OK.
   d. Click **OK**.

   CLI equivalent

```
# alerts notify-list add eng_lab emails mlee@urcompany.com,bob@urcompany.com
# alerts notify-list del eng_lab emails bob@urcompany.com
```

5. Click **OK**.

Deleting a notification group

Use the Notification tab to delete one or more existing notification groups.

Steps

1. Select **Health** > **Alerts** > **Notification**.
2. Select one or more checkboxes of groups in the Notifications group list, and click **Delete**.

The Delete Group dialog box appears.

3. Verify the deletion and click **OK**.

   CLI equivalent

```
# alerts notify-list destroy eng_grp
```

Resetting the notification group configuration

Use the Notification tab to remove all notification groups added and to remove any changes made to the Default group.

Steps

1. Select **Health** > **Alerts** > **Notification**.
2. Select **More Tasks** > **Reset Notification Groups**.
3. In the Reset Notification Groups dialog box, click **Yes** in the verification dialog.

   CLI equivalent

```
# alerts notify-list reset
```

Configuring the daily summary schedule and distribution list

Every day, each managed system sends a Daily Alert Summary email to the subscribers configured for the alerts summary.list email group. The Daily Alert Summary email contains current and historical alerts showing messages about non-critical hardware situations and disk space usage numbers that you might want to address soon.

About this task

A fan failure is an example of a noncritical issue that you might want to address as soon as is reasonably possible. When Support receives the failure notification, they contact you to arrange for component replacement.

Steps

1. Select **Maintenance** > **Support** > **Telemetry**.
2. Click **Configure**.
3. Specify the alert summary configuration:
   a. Select **Enable**.
   b. Select **Daily** or **Weekly**.
   c. Use the list boxes to select the hour, minute, and either AM or PM for the summary report.
   d. Select **Text** or **HTML**.
   e. Click **Next**.

      CLI equivalent

```
# autosupport set schedule alert-summary daily 1400
# autosupport set alert-summary-format {text | html}
```

4. Click through the rest of the Telemetry configuration screen.
5. Click **Close**.
6. To edit the configuration, click **Configure** and repeat steps 3-5.
7. To configure the daily alert subscriber list:.
   a. In the Alert Summary Subscribers panel, click **Add**.
   b. Specify an email address.
   c. Click **Add**.

      CLI equivalent

```
# autosupport add alert-summary emails djones@company.com
```

- To delete an email address, click **Remove** next to the email address.

  CLI equivalent

  ```
  # autosupport del alert-summary emails djones@company.com
  ```

8

Managing System Access

This chapter presents the following topics:

Topics:

System access management

System access management features allow you to control system access to users in a local database or in a network directory. Additional controls define different access levels and control which protocols can access the system.

Role-based access control

Role-based access control (RBAC) is an authorization policy that controls which DD System Manager controls and CLI commands a user can access on a system.

For example, users who are assigned the *admin* role can configure and monitor an entire system, while users who are assigned the *user* role are limited to monitoring a system. When logged into DD System Manager, users see only the program features that they are permitted to use based on the role assigned to the user. The following roles are available for administering and managing the DDOS.

admin	An *admin* role user can configure and monitor the entire system. Most configuration features and commands are available only to *admin* role users. However, some features and commands require the approval of a *security* role user before a task is completed.
limited-admin	The *limited-admin* role can configure and monitor the system with some limitations. Users who are assigned this role cannot perform data deletion operations, edit the registry, or enter bash or SE mode.
user	The *user* role enables users to monitor systems and change their own password. Users who are assigned the *user* management role can view system status, but they cannot change the system configuration.
security (security officer)	A *security* role user, who may be referred to as a security officer, can manage other security officers, authorize procedures that require security officer approval, and perform all tasks supported for user-role users. The *security* role is provided to comply with the Write Once Read-Many (WORM) regulation. This regulation requires electronically stored corporate data be kept in an unaltered, original state for purposes such as eDiscovery, auditing, and logging. As a result of compliance regulations, most command options for administering sensitive operations, such as DD Encryption, DD Retention Lock Compliance, and archiving now require security officer approval. In a typical scenario, an *admin* role user issues a command and, if security officer approval is required, the system displays a prompt for approval. To proceed with the original task, the security officer must enter his or her username and password on the same console at which the command was run. If the system recognizes the security officer credentials, the procedure is authorized. If not, a security alert is generated. The following are some guidelines that apply to security-role users: ● Only the *sysadmin* user (the default user created during the DDOS installation) can create the first security officer, after which the privilege to create security officers is removed from the *sysadmin* user. **NOTE:** The *sysadmin* user cannot set the first security officer password to the same value as the *sysadmin* password.

- The system prompts to create a security officer the first time the *sysadmin* logs into the system via SSH or Telnet. If the *sysadmin* user chooses to create the security officer at that time, they must select **yes** and specify the username and password for the security officer.
- After the first security officer is created, only security officers can create other security officers.
- Creating a security officer does not enable the authorization policy. To enable the authorization policy, a security officer must log in and enable the authorization policy.
- Separation of privilege and duty apply. *admin* role users cannot perform security officer tasks, and security officers cannot perform system configuration tasks.
- During an upgrade, if the system configuration contains security officers, a sec-off-defaults permission is created that includes a list of all current security officers.

backup-operator	A *backup-operator* role user can perform all tasks permitted for *user* role users, create snapshots for MTrees, import, export, and move tapes between elements in a virtual tape library, and copy tapes across pools. A *backup-operator* role user can also add and delete SSH public keys for non-password-required log ins. (This function is used mostly for automated scripting.) He or she can add, delete, reset and view CLI command aliases, synchronize modified files, and wait for replication to complete on the destination system.
none	The *none* role is for DD Boost authentication and tenant-unit users only. A *none* role user can log in to a protection system and can change his or her password, but cannot monitor, manage, or configure the primary system. When the primary system is partitioned into tenant units, either the *tenant-admin* or the *tenant-user* role is used to define a user's role with respect to a specific tenant unit. The tenant user is first assigned the *none* role to minimize access to the primary system, and then either the *tenant-admin* or the *tenant-user* role is appended to that user.
tenant-admin	A *tenant-admin* role can be appended to the other (non-tenant) roles when the Secure Multi-Tenancy (SMT) feature is enabled. A *tenant-admin* user can configure and monitor a specific tenant unit.
tenant-user	A *tenant-user* role can be appended to the other (non-tenant) roles when the SMT feature is enabled. The *tenant-user* role enables a user to monitor a specific tenant unit and change the user password. Users who are assigned the *tenant-user* management role can view tenant unit status, but they cannot change the tenant unit configuration.

Access management for IP protocols

This feature manages system access for the FTP, FTPS, HTTP, HTTPS, SSH, SCP, and Telnet protocols.

Viewing the IP services configuration

The Administrator Access tab displays the configuration status for the IP protocols that can be used to access the system. FTP and FTPS are the only protocols that are restricted to administrators.

Steps

Select **Administration** > **Access** > **Administrator Access**.

Results

The Access Management page displays the Administrator Access, Local Users, Authentication, and Active Users tabs.

Managing FTP access

The File Transfer Protocol (FTP) allows administrators to access files on the protection system.

About this task

You can enable either FTP or FTPS access to users who are assigned the admin management role. FTP access allows admin user names and passwords to cross the network in clear text, making FTP an insecure access method. FTPS is recommended as a secure access method. When you enable either FTP or FTPS access, the other access method is disabled.

NOTE: Only users who are assigned the admin management role are permitted to access the system using FTP

NOTE: LFTP clients that connect to a protection system via FTPS or FTP are disconnected after reaching a set timeout limit. However the LFTP client uses its cached username and password to reconnect after the timeout while you are running any command.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. Select **FTP** and click **Configure**.
3. To manage FTP access and which hosts can connect, select the General tab and do the following:
   a. To enable FTP access, select **Allow FTP Access**.
   b. To enable all hosts to connect, select **Allow all hosts to connect**.
   c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the Allowed Hosts list.

      NOTE: You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

      - To add a host, click Add (**+**). Enter the host identification and click **OK**.
      - To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
      - To remove a host ID, select the host in the **Hosts** list and click Delete (**X**).
4. To set a session timeout, select the **Advanced** tab, and enter the timeout value in seconds.

   NOTE: The session timeout default is Infinite, that is, the connection does not close.

5. Click **OK**.
   If FTPS is enabled, a warning message appears with a prompt to click **OK** to proceed.

Managing FTPS access

The FTP Secure (FTPS) protocol allows administrators to access files on the protection system.

About this task

FTPS provides additional security over using FTP, such as support for the Transport Layer Security (TLS) and for the Secure Sockets Layer (SSL) cryptographic protocols. Consider the following guidelines when using FTPS.

- Only users who are assigned the admin management role are permitted to access the system using FTPS.
- When you enable FTPS access, FTP access is disabled.
- When you issue the `get` command, the fatal error message `SSL_read: wrong version number lftp` appears if matching versions of SSL are not installed on the protection system and compiled on the LFTP client . As a workaround, attempt to re-issue the `get` command on the same file.

By default, FTPS is configured to use:

- TLS version 1.2

  NOTE: Run the `adminaccess ftps option set tls-version` command to change the TLS version.
- SSL/TLS cipher list "ALL:!ADH:!EXPORT56:!EXPORT40:+HIGH:!MEDIUM:!LOW:!SSLv2:!SSLv3:!DES-CBC3-SHA:+EXP@STRENGTH"

  NOTE: Run the `adminaccess ftps option set cipher-list` command to change the cipher list.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. Select **FTPS** and click **Configure**.
3. To manage FTPS access and which hosts can connect, select the General tab and do the following:
   a. To enable FTPS access, select **Allow FTPS Access**.
   b. To enable all hosts to connect, select **Allow all hosts to connect**.
   c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the hosts list.

      NOTE: You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

- To add a host, click Add (**+**). Enter the host identification and click **OK**.
- To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
- To remove a host ID, select the host in the **Hosts** list and click Delete (**X**).

4. To set a session timeout, select the **Advanced** tab and enter the timeout value in seconds.

   NOTE: The session timeout default is Infinite, that is, the connection does not close.

5. Click **OK**. If FTP is enabled, a warning message appears and prompts you to click **OK** to proceed.

Managing HTTP and HTTPS access

HTTP or HTTPS access is required to support browser access to DD System Manager. By default, GUI access over HTTPS is enabled and GUI access over HTTP is disabled.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. Select **HTTP** or **HTTPS** and click **Configure**.

   The Configure HTTP/HTTPS Access dialog appears and displays tabs for general configuration, advanced configuration, and certificate management.
3. To manage the access method and which hosts can connect, select the General tab and do the following:
   a. Select the checkboxes for the access methods you want to allow.
   b. To enable all hosts to connect, select **Allow all hosts to connect**.
   c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

      NOTE: You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

      - To add a host, click Add (**+**). Enter the host identification and click **OK**.
      - To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
      - To remove a host ID, select the host in the **Hosts** list and click Delete (**X**).
4. To configure system ports and session timeout values, select the **Advanced** tab, and complete the form.
   - In the **HTTP Port** box, enter the port number. Port 80 is assigned by default.
   - In the **HTTPS Port** box, enter the number. Port 443 is assigned by default.
   - In the **Session Timeout** box, enter the interval in seconds that must elapse before a connection closes. The minimum is 60 seconds and the maximum is 31536000 seconds (one year).

   NOTE: The session timeout default is 10,800 seconds.

5. Click **OK**.

Managing host certificates for HTTP and HTTPS

A host certificate allows browsers to verify the identity of the system when establishing management sessions.

Requesting a host certificate for HTTP and HTTPS

You can use DD System Manager to generate a host certificate request, which you can then forward to a Certificate Authority (CA).

About this task

NOTE: You must configure a system passphrase (system passphrase set) before you can generate a CSR.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.

3. Select the **Certificate** tab.
4. Click Add.

   A dialog appears for the protocol you selected earlier in this procedure.

5. Click **Generate the CSR for this Data Domain system**.
   The dialog expands to display a CSR form.

   NOTE: DDOS supports one active CSR at a time. After a CSR is generated, the **Generate the CSR for this Data Domain system** link is replaced with the **Download the CSR for this Data Domain system** link. To delete a CSR, use the `adminaccess certificate cert-signing-request delete` CLI command.

6. Complete the CSR form and click **Generate and download a CSR**.

   The CSR file is saved at the following path: `/ddvar/certificates/CertificateSigningRequest.csr`. Use SCP, FTP or FTPS to transfer the CSR file from the system to a computer from which you can send the CSR to a CA.

Adding a host certificate for HTTP and HTTPS

You can use DD System Manager to add a host certificate to the system.

Steps

1. If you did not requested a host certificate, request a host certificate from a certificate authority.
2. When you receive a host certificate, copy or move it to the computer from which you run DD Service Manager.
3. Select **Administration** > **Access** > **Administrator Access**.
4. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.
5. Select the **Certificate** tab.
6. Click **Add**.

   A dialog appears for the protocol you selected earlier in this procedure.

7. To add a host certificate enclosed in a .p12 file, do the following:
   a. Select **I want to upload the certificate as a .p12 file**.
   b. Type the password in the **Password** box.
   c. Click **Browse** and select the host certificate file to upload to the system.
   d. Click **Add.**
8. To add a host certificate enclosed in a .pem file, do the following:
   a. Select **I want to upload the public key as a .pem file and use a generated private key**.
   b. Click **Browse** and select the host certificate file to upload to the system.
   c. Click **Add.**

Deleting a host certificate for HTTP and HTTPS

DDOS supports one host certificate for HTTP and HTTPS. If the system is currently using a host certificate and you want to use a different host certificate, you must delete the current certificate before adding the new certificate.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. In the Services area, select **HTTP** or **HTTPS** and click **Configure**.
3. Select the **Certificate** tab.
4. Select the certificate you want to delete.
5. Click **Delete**, and click **OK**.

Managing SSH and SCP access

SSH is a secure protocol that enables network access to the system CLI, with or without SCP (secure copy). You can use DD System Manager to enable system access using the SSH protocol. SCP requires SSH, so when SSH is disabled, SCP is automatically disabled.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. Select **SSH** or **SCP** and click **Configure**.
3. To manage the access method and which hosts can connect, select the **General** tab.
   a. Select the checkboxes for the access methods you want to allow.
   b. To enable all hosts to connect, select **Allow all hosts to connect**.
   c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

      NOTE: You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

      - To add a host, click Add (**+**). Enter the host identification and click **OK**.
      - To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
      - To remove a host ID, select the host in the **Hosts** list and click Delete (**X**).
4. To configure system ports and session timeout values, click the **Advanced** tab.
   - In the **SSH/SCP** Port text entry box, enter the port number. Port 22 is assigned by default.
   - In the **Session Timeout** box, enter the interval in seconds that must elapse before connection closes.

   NOTE: The session timeout default is Infinite, that is, the connection does not close.

   NOTE: Click **Default** to revert to the default value.
5. Click **OK**.

Managing Telnet access

Telnet is an insecure protocol that enables network access to the system CLI.

About this task

NOTE: Telnet access allows user names and passwords to cross the network in clear text, making Telnet an insecure access method.

Steps

1. Select **Administration** > **Access** > **Administrator Access**.
2. Select **Telnet** and click **Configure**.
3. To manage Telnet access and which hosts can connect, select the **General** tab.
   a. To enable Telnet access, select **Allow Telnet Access**.
   b. To enable all hosts to connect, select **Allow all hosts to connect**.
   c. To restrict access to select hosts, select **Limit Access to the following systems**, and modify the host list.

      NOTE: You can identify a host using a fully qualified hostname, an IPv4 address, or an IPv6 address.

      - To add a host, click Add (**+**). Enter the host identification and click **OK**.
      - To modify a host ID, select the host in the **Hosts** list and click Edit (pencil). Change the host ID and click **OK**.
      - To remove a host ID, select the host in the **Hosts** list and click Delete (**X**).
4. To set a session timeout, select the Advanced tab and enter the timeout value in seconds.

   NOTE: The session timeout default is Infinite, that is, the connection does not close.
5. Click **OK**.

Local user account management

A local user is a user account (user name and password) that is configured on the protection system instead of being defined in a Windows Active Directory, Windows Workgroup, or NIS directory.

After a trusted domain is configured, users who belong to that domain will be able to log into the protection system even if that trusted domain is offline.

UID conflicts: local user and NIS user accounts

When you set up a protection system in an NIS environment, be aware of potential UID conflicts between local and NIS user accounts.

Local user accounts on a protection system start with a UID of 500. To avoid conflicts, do not use UIDs less than 1000 to avoid conflict with local user accounts on the system, and consider the size of potential local accounts when you define allowable UID ranges for NIS users.

Viewing local user information

Local users are user accounts that are defined on the system, rather than in Active Directory, a Workgroup, or UNIX. You can display the local user's username, management role, login status, and target disable date. You can also display the user's password controls and the tenant units the user can access.

About this task

NOTE: The user-authentication module uses Greenwich Mean Time (GMT). To ensure that user accounts and passwords expire correctly, configure settings to use the GMT that corresponds to the target local time.

Steps

1. Select **Administration** > **Access** > **Local Users** .

   The Local Users view appears and shows the Local Users table and the Detailed Information area.

   Table 4. Local user list column label descriptions

Item	Description
Name	The user ID, as added to the system.
Management Role	The role displayed is admin, user, security, backup-operator, or none. In this table, Tenant user roles are displayed as *none*. To see an assigned tenant role, select the user and view the role in the Detailed Information area.
Status	• Enabled: User access to the account is permitted. • Disabled: User access to the account is disabled and requires a system administrator to enable. • Password-aged: User password is aged and expired. The user will be prompted to confirm their current password and create a new password on the next login.
Disable Date	The date the account is set to be disabled.
Last Login From	The location where the user last logged in.
Last Login Time	The time the user last logged in.

   NOTE: User accounts configured with the admin or security officer roles can view all users. Users with other roles can view only their own user accounts.

2. Select the user you want to view from the list of users.

   Information about the selected user displays in the Detailed Information area.

   NOTE: The default values are the initial default password policy values. A system administrator (admin role) can change them by selecting **More Tasks** > **Change Login Options**.

Creating local users

Create local users when you want to manage access on the local system instead of through an external directory. Protection systems support a maximum of 500 local user accounts.

Steps

1. Select **Administration** > **Access** > **Local Users**.

   The Local Users view appears.

2. Click **Create** to create a new user.

   The Create User dialog appears.

3. Enter user information in the General Tab.

   The default value for the minimum length of a password is 9 characters. The default value for the minimum number of character classes required for a user password is 4. Allowable character classes include:

   - Lowercase letters (a-z)
   - Uppercase letters (A-Z)
   - Numbers (0-9)
   - Special Characters ($, %, #, +, and so on)

   NOTE: Sysadmin is the default admin-role user and cannot be deleted or modified.

4. To manage password and account expiration, select the Advanced tab and set the controls as required.
5. Click **OK**.

   NOTE: Note: The default password policy can change if an admin-role user changes them (**More Tasks** > **Change Login Options**). The default values are the initial default password policy values.

Modifying a local user profile

After you create a user, you can use DD System Manager to modify the user configuration.

Steps

1. Select **Administration** > **Access** > **Local Users**.

   The Local Users view appears.

2. Click a user name from the list.
3. Click **Modify** to make changes to a user account.

   The Modify User dialog box appears.

4. Update the information on the General tab.

   NOTE: If SMT is enabled and a role change is requested from none to any other role, the change is accepted only if the user is not assigned to a tenant-unit as a management-user, is not a DD Boost user with its default-tenant-unit set, and is not the owner of a storage-unit that is assigned to a tenant-unit.

   NOTE: To change the role for a DD Boost user that does not own any storage units, unassign it as a DD Boost user, change the user role, and re- assign it as a DD Boost user again.

5. Update the information on the Advanced tab.
6. Click **OK**.

Deleting a local user

You can delete certain users based on your user role. If one of the selected users cannot be deleted, the Delete button is disabled.

About this task

The sysadmin user cannot be deleted. Admin users cannot delete security officers. Only security officers can delete, enable, and disable other security officers.

Steps

1. Select **Administration** > **Access** > **Local Users**.

   The Local Users view appears.

2. Click one or more user names from the list.
3. Click **Delete** to delete the user accounts.

   The Delete User dialog box appears.

4. Click **OK** and **Close**.

Enabling and disabling local users

Admin users can enable or disable all users except the sysadmin user and users with the security role. The sysadmin user cannot be disabled. Only Security officers can enable or disable other security officers.

Steps

1. Select **Administration** > **Access** > **Local Users**.

   The Local Users view appears.

2. Click one or more user names from the list.
3. Click either **Enable** or **Disable** to enable or disable user accounts.

   The Enable or Disable User dialog box appears.

4. Click **OK** and **Close**.

Enabling security authorization

You can use the CLI to enable and disable the security authorization policy.

About this task

For information on the commands used in this procedure, see the *DDOS Command Reference Guide*.

NOTE: The DD Retention Lock Compliance license must be installed. You are not permitted to disable the authorization policy on DD Retention Lock Compliance systems.

Steps

1. Log into the CLI using a security officer username and password.
2. To enable the security officer authorization policy, enter: # `authorization policy set security-officer enabled`

Changing user passwords

After you create a user, you can use DD System Manager to change the user's password. Individual users can also change their own passwords by selecting **User Options** > **Change Password** from the top banner of DD System Manager .

Steps

1. Click **Administration** > **Access** > **Local Users**.

   The Local Users view is displayed.

2. Click a username from the list.
3. To change the user password, click **Change Password**.

   The Change Password dialog box is displayed.

4. Enter the old password into the **Old Password** box.
5. Enter the new password into the **New Password** box.
6. Enter the new password again into **Verify New Password** box.
7. Click **OK**.

   Only users with an "admin" role may change the password of other users. The administrator can change the password of other users from the CLI by running the `user change password [<user>]` command.

   NOTE: For security reasons, users with an "admin" role cannot change other "admin" users' passwords. If an "admin" user password needs to be changed by logging in as another user, contact Dell Support by creating a Support Request or chat request for assistance.

Modifying the password policy and login controls

The password policy and login controls define login requirements for all users. Administrators can specify how often a password must be changed, what is required to create a valid password, and how the system responds to invalid login attempts.

Steps

1. Select **Administration** > **Access**.
2. Select **More Tasks** > **Change Login Options**.

   The Change Login Options dialog appears.

3. Specify the new configuration in the boxes for each option. To select the default value, click **Default** next to the appropriate option.
4. Click **OK** to save the password settings.

Change Login Options dialog

Use this dialog to set the password policy and specify the maximum login attempts and lockout period.

Table 5. Change Login Options dialog controls

Item	Description
Minimum Days Between Change	The minimum number of days between password changes that you allow a user. This value must be less than the **Maximum Days Between Change** value minus the **Warn Days Before Expire** value. The default setting is 0.
Maximum Days Between Change	The maximum number of days between password changes that you allow a user. The minimum value is 1. The default value is 99999.
Warn Days Before Expire	The number of days to warn the users before their password expires. This value must be less than the **Maximum Days Between Change** value minus the **Minimum Days Between Change** value. The default setting is 7.

Table 5. Change Login Options dialog controls (continued)

Item	Description
Disable Days After Expire	The system disables a user account after password expiration according to the number of days specified with this option. Valid entries are *never* or number greater than or equal to 0. The default setting is never.
Minimum Length of Password	The minimum password length required. The allowed range is from 9 to 31. Default is 9.
Maximum Three Consecutive Characters	Enable or disable the requirement for a maximum of three repeated characters. The default setting is enabled.
Number of Previous Passwords to Block	Specify the number of remembered passwords. The range is 0 to 24, and the default settings is 6. **NOTE:** If this setting is reduced, the remembered password list remains unchanged until the next time the password is changed. For example, if this setting is changed from 4 to 3, the last four passwords are remembered until the next time the password is changed.
Maximum login attempts	Specifies the maximum number of login attempts before a mandatory lock is applied to a user account. This limit applies to all user accounts, including sysadmin. A locked user cannot log in while the account is locked. The range is 3 to 20, and the default value is 3.
Unlock timeout (seconds)	Specifies how long a user account is locked after the maximum number of login attempts. When the configured unlock timeout is reached, a user can attempt login. The range is 120 to 3600 seconds, and the default period is 120 seconds.
Minimum positions changed	Specify the minimum number of positions that must be changed when a password is changed. The default value is 1.
Maximum active logins	Specifies the maximum number of active logins to allow. The default value is 100.

Directory user and group management

You can use DD System Manager to manage access to the system for users and groups in Windows Active Directory, Windows Workgroup, and NIS. Kerberos authentication is an option for CIFS and NFS clients.

Viewing Active Directory and Kerberos information

The Active Directory Kerberos configuration determines the methods CIFS and NFS clients use to authenticate. The Active Directory/Kerberos Authentication panel displays this configuration.

Steps

1. Select **Administration** > **Access** > **Authentication**.
2. Expand the Active Directory/Kerberos Authentication panel.

Configuring Active Directory and Kerberos authentication

Configuring Active Directory authentication makes the protection system part of a Windows Active Directory realm. CIFS clients and NFS clients use Kerberos authentication.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.
   The Active Directory/Kerberos Authentication dialog appears.
4. Select **Windows/Active Directory** and click **Next**.

5. Enter the full realm name for the system (for example: domain1.local), the user name, and password for the system. Then click **Next**.

   NOTE: Use the complete realm name. Ensure that the user is assigned sufficient privileges to join the system to the domain. The user name and password must be compatible with Microsoft requirements for the Active Directory domain. This user must also be assigned permission to create accounts in this domain.

6. Select the default CIFS server name, or select **Manual** and enter a CIFS server name.
7. To select domain controllers, select **Automatically assign**, or select **Manual** and enter up to three domain controller names. You can enter fully qualified domain names, hostnames, or IP (IPv4 or IPv6) addresses.
8. To select an organizational unit, select **Use default Computers**, or select **Manual** and enter an organization unit name.

   NOTE: The account is moved to the new organizational unit.

9. Click **Next**.
   The Summary page for the configuration appears.
10. Click **Finish**.
    The system displays the configuration information in the Authentication view.
11. To enable administrative access, click **Enable** to the right of **Active Directory Administrative Access**.

Authentication mode selections

The authentication mode selection determines how CIFS and NFS clients authenticate using supported combinations of Active Directory, Workgroup, and Kerberos authentication.

About this task

DDOS supports the following authentication options.

- Disabled: Kerberos authentication is disabled for CIFS and NFS clients. CIFS clients use Workgroup authentication.
- Windows/Active Directory: Kerberos authentication is enabled for CIFS and NFS clients. CIFS clients use Active Directory authentication.
- Unix: Kerberos authentication is enabled for only NFS clients. CIFS clients use Workgroup authentication.

Managing administrative groups for Active Directory

You can use the Active Directory/Kerberos Authentication panel to create, modify, and delete Active Directory (Windows) groups and assign management roles (admin, backup-operator, and so on) to those groups.

To prepare for managing groups, select **Administration** > **Access** > **Authentication** , expand the Active Directory/Kerberos Authentication panel, and click the Active Directory Administrative Access **Enable** button.

Creating administrative groups for Active Directory

Create an administrative group when you want to assign a management role to all the users configured in an Active Directory group.

Prerequisites

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration** > **Access** > **Authentication** page.

Steps

1. Click **Create...**.
2. Enter the domain and group name separated by a backslash. For example: domainname\groupname.
3. Select the management role for the group from the drop-down menu.
4. Click **OK**.

Modifying administrative groups for Active Directory

Modify an administrative group when you want to change the administrative domain name or group name configured for an Active Directory group.

Prerequisites

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration** > **Access** > **Authentication** page.

Steps

1. Select a group to modify under the **Active Directory Administrative Access** heading.
2. Click **Modify...**.
3. Modify the domain and group name. These names are separated by a backslash. For example: domainname\groupname.

Deleting administrative groups for Active Directory

Delete an administrative group when you want to terminate system access for all the users configured in an Active Directory group.

Prerequisites

Enable Active Directory Administrative Access on the Active Directory/Kerberos Authentication panel in the **Administration** > **Access** > **Authentication** page.

Steps

1. Select a group to delete under the **Active Directory Administrative Access** heading.
2. Click **Delete**.

Configuring UNIX Kerberos authentication

Configuring UNIX Kerberos authentication enables NFS clients to use Kerberos authentication. CIFS clients use Workgroup authentication.

Prerequisites

NIS must be running for UNIX-mode Kerberos authentication to function. For instructions about enabling Kerberos, see the section regarding enabling NIS services.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.
   The Active Directory/Kerberos Authentication dialog appears.
4. Select **Unix** and click **Next**.
5. Enter the realm name (for example: domain1.local), and up to three host names or IP addresses (IPv4 or IPv6) for key distribution centers (KDCs).
6. Optionally, click **Browse** to upload a keytab file, and click **Next**.
   The Summary page for the configuration appears.

   NOTE: Keytab files are generated on the authentication servers (KDCs) and contain a shared secret between the KDC server and the DDR.

   NOTE: A keytab file must be uploaded and imported for Kerberos authentication to operate correctly.

7. Click **Finish**.
   The system displays the configuration information in the Active Directory/Kerberos Authentication panel.

Disabling Kerberos authentication

Disabling Kerberos authentication prevents CIFS and NFS clients from using Kerberos authentication. CIFS clients use Workgroup authentication.

Steps

1. Select **Administration** > **Access Management** > **Authentication**.

   The Authentication view appears.

2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Configure...** next to Mode to start the configuration wizard.
   The Active Directory/Kerberos Authentication dialog appears.
4. Select **Disabled** and click **Next**.
   The system displays a summary page with changes appearing in bold text.
5. Click **Finish**.
   The system displays Disabled next to Mode in the Active Directory/Kerberos Authentication panel.

Viewing Workgroup authentication information

Use the Workgroup Authentication panel to view Workgroup configuration information.

Steps

1. Select **Administration** > **Access** > **Authentication**.
2. Expand the Workgroup Authentication panel.

Configuring workgroup authentication parameters

Workgroup authentication parameters allow you to configure a Workgroup name and CIFS server name.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Workgroup Authentication panel.
3. Click **Configure**.

   The Workgroup Authentication dialog appears.

4. For Workgroup Name, select **Manual** and enter a workgroup name to join, or use the default.

   The Workgroup mode joins a protection system to a workgroup domain.

5. For CIFS Server Name, select **Manual** and enter a server name (the DDR), or use the default.
6. Click **OK**.

Viewing LDAP authentication information

The LDAP Authentication panel displays the LDAP configuration parameters and whether LDAP authentication is enabled or disabled.

About this task

Enabling LDAP allows you to use an existing OpenLDAP server or deployment with the protection system for system-level user authentication, NFSv4 ID mapping, NFSv3 Kerberos with LDAP, or NFSv4 Kerberos with LDAP.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the LDAP Authentication panel.

Enabling and disabling LDAP authentication

Use the LDAP authentication panel to enable, disable, or reset LDAP authentication.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the LDAP authentication panel.
3. Click **Enable** next to LDAP Status to enable or **Disable** to disable LDAP Authentication.

   The Enable or Disable LDAP authentication dialog box appears.

   NOTE: An LDAP server must exist before enabling LDAP authentication.

4. Click **OK**.

Resetting LDAP authentication.

The **Reset** button disables LDAP authentication and clears the LDAP configuration information.

Configuring LDAP authentication

Use the LDAP authentication panel to configure LDAP authentication.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the LDAP Authentication panel.
3. Click **Configure**.

   The Configure LDAP Authentication dialog box appears.

4. Specify the base suffix in the **Base Suffix** field.
5. Specify the account name to associate with the LDAP server in the **Bind DN** field.
6. Specify the password for the Bind DN account in the **Bind Password** field.
7. Optionally select **Enable SSL**.
8. Optionally select **Demand server certificate** to require the protection system to import a CA certificate from the LDAP server.
9. Click **OK**.
10. If necessary at a later time, click **Reset** to return the LDAP configuration to its default values.

Specifying LDAP authentication servers

Use the LDAP authentication panel to specify LDAP authentication servers.

Prerequisites

LDAP authentication must be disabled before configuring an LDAP server.

About this task

NOTE: DD SM performance when logging in with LDAP will decrease as the number of hops between the system and the LDAP server increaes.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the LDAP authentication panel.
3. Click the **+** button to add a server.
4. Specify the LDAP server in one of the following formats:
   - IPv4 address—**`10.26.16.250`**
   - IPv6 address—**`[::ffff:9.53.96.21]`**
   - Hostname—**`myldapserver`**
5. Click **OK**.

Configuring LDAP groups

Use the LDAP authentication panel to configure LDAP groups.

About this task

LDAP group configuration only applies when using LDAP for user authentication on the protection system.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the LDAP authentication panel.
3. Configure the LDAP groups in the LDAP Group table.
   - To add an LDAP group, click Add (**+**), enter the LDAP group name and role, and click**OK**.
   - To modify an LDAP group, select the checkbox of the group name in the LDAP group list and click Edit (pencil). Change the LDAP group name and click **OK**.
   - To remove an LDAP group, select the LDAP group in the list and click Delete (**X**).

Using the CLI to configure LDAP authentication

You can use the CLI to configure an existing OpenLDAP server or deployment with a protection system for system-level user authentication, NFSv4 ID mapping, NFSv3 Kerberos with LDAP, or NFSv4 Kerberos with LDAP.

This cannot be configured if LDAP authentication is already configured for Active Directory.

Configuring LDAP authentication for Active Directory

DDOS supports the use of LDAP authentication for Active Directory. When using LDAP authentication with Active Directory, CIFS data access for Active Directory users and groups is not allowed, and CIFS shares on the system are accessible only to local users. Only CLI and GUI logins are allowed for Active Directory users with this configuration.

Prerequisites

Make sure the environment meets the following requirements to configure LDAP authentication for Active Directory:

- TLS/SSL is enabled for LDAP communication.
- Active Directory users accessing the protection system must have valid UID and GID numbers.
- Active Directory groups accessing the protection system must have a valid GID number.

NOTE:

- Specify the username in the format *<username>*, without specifying a domain name.
- Specify the groupame in the format *<groupname>*, without specifying a domain name.

  NOTE: Groupnames that use a dot (.) are not supported for authentication.
- User and group names are not case sensitive.

The following limitations apply to LDAP for Active Directory:

- Microsoft Active Directory is the only supported Active Directory provider.
- Active Directory LDS (Lightweight Directory Services) is not supported.
- The Active Directory native schema for uidNumber and gidNumber population is the only supported schema. There is no support for third-party tools integrated with Active Directory.

About this task

LDAP authentication for Active Directory cannot be used in combination with Active Directory or Kerberos authentication for CIFS. The CLI is the only way to configure this option.

Steps

Run the authentication **`ldap base set basename type active-directory`** command to enable LDAP authentication for Active Directory.

NOTE: The command will fail if the CIFS authentication is already configured as Active Directory.

```
# authentication ldap base set "dc=anvil,dc=team" type active-directory
```

Configure LDAP servers

You can configure one or more LDAP servers at the same time. Configure servers from the site nearest to the protection system for minimum latency.

About this task

NOTE: LDAP must be disabled when making any changes to the configuration.

Specify the LDAP server in one of the following formats:

- IPv4 address—**`10.<A>.<B>.<C>`**
- IPv4 address with port number—**`10.<A>.<B>.<C>:400`**
- IPv6 address—**`[::ffff:9.53.96.21]`**
- IPv6 address with port number—**`[::ffff:9.53.96.21]:400`**
- Hostname—**`myldapserver`**
- Hostname with port number—**`myldapserver:400`**

When configuring multiple servers:

- Separate each server with a space.
- The first server listed when using the `authentication ldap servers add` command becomes the primary server.
- If any of the servers cannot be configured, the command fails for all servers listed.

Steps

1. Add one or more LDAP servers by using the `authentication ldap servers add` command:

```
# authentication ldap servers add 10.A.B.C 10.X.Y.Z:400
LDAP server(s) added
LDAP Server(s):        2
#     IP Address/Hostname
---   --------------------
1.    10.A.B.C (primary)
2.    10.X.Y.Z:400
---   --------------------
```

2. Remove one or more LDAP servers by using the `authentication ldap servers del` command:

```
# authentication ldap servers del 10.X.Y.Z:400
LDAP server(s) deleted.
LDAP Servers: 1
#   Server
-   ------------   ---------
1   10.A.B.C       (primary)
-   ------------   ---------
```

3. Remove all LDAP servers by using the `authentication ldap servers reset` command:

```
# authentication ldap servers reset
LDAP server list reset to empty.
```

Configure the LDAP base suffix

The base suffix is the base DN for search and is where the LDAP directory begins searching.

About this task

Set the base suffix for OpenLDAP or Active Directory.

NOTE: The base suffix cannot be set for both OpenLDAP and Active Directory

User login is allowed from the primary Active Directory domain only. Users and groups from trusted Active Directory domains are not supported.

Set the base suffix for OpenLDAP

Steps

Set the LDAP base suffix by using the `authentication ldap base set` command:

```
# authentication ldap base set "dc=anvil,dc=team"
LDAP base-suffix set to "dc=anvil,dc=team".
```

Set the base suffix for Active Directory

Steps

1. Set the LDAP base suffix by using the `authentication ldap base set` command:

```
# authentication ldap base set "dc=anvil,dc=team" type active-directory
LDAP base-suffix set to "dc=anvil,dc=team".
```

2. Set the LDAP group by using the `authentication ldap groups add` command:

   NOTE: In this example, all users in the dd-admins LDAP group will have administrative privileges on the protection system.

```
# authentication ldap groups add dd-admins role admin
LDAP Group   Role
----------   -----
dd-admins    admin
----------   -----
```

Reset the LDAP base suffix

Steps

Reset the LDAP base suffix by using the `authentication ldap base reset` command:

```
# authentication ldap base reset
LDAP base-suffix reset to empty.
```

Configure LDAP client authentication

Configure the account (Bind DN) and password (Bind PW) that is used to authenticate with the LDAP server and make queries.

About this task

You should always configure the Bind DN and password. Normally, LDAP servers require authenticated bind by default. If `client-auth` is not set, anonymous access is requested, providing no name or password. The output of `authentication ldap show` is as follows:

```
# authentication ldap show
LDAP configuration
        Enabled:          yes (*)
        Base-suffix:      dc=u2,dc=team
        Binddn:           (anonymous)
        Server(s):        1
#   Server
-   -------------   ---------
1   10.207.86.160   (primary)
-   -------------   ---------

Secure LDAP configuration
        SSL Enabled:      no
        SSL Method:       off
        tls_reqcert:      demand

(*) Requires a filesystem restart for the configuration to take effect.
```

If binddn is set using `client-auth` CLI, but bindpw is not provided, unauthenticated access is requested.

```
# authentication ldap client-auth set binddn "cn=Manager,dc=u2,dc=team"
Enter bindpw:
** Bindpw is not provided. Unauthenticated access would be requested.
LDAP client authentication binddn set to "cn=Manager,dc=u2,dc=team".
```

Steps

1. Set the Bind DN and password by using the `authentication ldap client-auth set binddn` command:

   ```
   # authentication ldap client-auth set binddn "cn=Administrator,cn=Users,dc=anvil,dc=team"
       Enter bindpw:
       LDAP client authentication binddn set to
   "cn=Administrator,cn=Users,dc=anvil,dc=team".
   ```

2. Reset the Bind DN and password by using the `authentication ldap client-auth reset` command:

   ```
   # authentication ldap client-auth reset
   LDAP client authentication configuration reset to empty.
   ```

Enable LDAP

Prerequisites

An LDAP configuration must exist before enabling LDAP. Additionally, you must disable NIS, ensure that the LDAP server is reachable, and be able to query the root DSE of the LDAP server.

Steps

1. Enable LDAP by using the `authentication ldap enable` command:

   ```
   # authentication ldap enable
   ```

The details of the LDAP configuration are displayed for you to confirm before continuing. To continue, type **yes** and restart the file system for LDAP configuration to take effect.

2. View the current LDAP configuration by using the `authentication ldap show` command:

   NOTE: If the system is configured to use LDAP for Active Directory, the command output will include a `Server Type` field to indicate it is connected to an Active Directory server.

```
# authentication ldap show
LDAP configuration
        Enabled:          no
        Base-suffix:      dc=anvil,dc=team
        Binddn:           cn=Administrator,cn=Users,dc=anvil,dc=team
        Server(s):        2
#   Server
-   ----------------   ---------
1   10.26.16.250       (primary)
2   10.26.16.251:400
-   ----------------   ---------

Secure LDAP configuration
        SSL Enabled:      no
        SSL Method:       off
        tls_reqcert:      demand
```

   Basic LDAP and secure LDAP configuration details are displayed.

3. View the current LDAP status by using the `authentication ldap status` command:

```
# authentication ldap status
```

   The LDAP status is displayed. If the LDAP status is not `good`, the problem is identified in the output. For example:

```
# authentication ldap status
Status: invalid credentials
```

   or

```
# authentication ldap status
Status: invalid DN syntax
```

4. Disable LDAP by using the `authentication ldap disable` command:

```
# authentication ldap disable
LDAP is disabled.
```

Enable secure LDAP

You can configure DDR to use secure LDAP by enabling SSL. For LDAP for Active Directory, configure secure LDAP with SSL/TLS options.

Prerequisites

If there is no LDAP CA certificate and `tls_reqcert` is set to `demand`, the operation fails. Import an LDAP CA certificate and try again.

If `tls_reqcert` is set to `never`, an LDAP CA certificate is not required. For more information, see Configure LDAP server certificate verification with imported CA certificates.

Steps

1. Enable SSL by using the `authentication ldap ssl enable` command:

```
# authentication ldap ssl enable
Secure LDAP is enabled with 'ldaps' method.
```

   The default method is secure LDAP, or *ldaps*. You can specify other methods, such as TLS:

```
# authentication ldap ssl enable method start_tls
Secure LDAP is enabled with 'start_tls' method.
```

2. Disable SSL by using the `authentication ldap ssl disable` command:

```
# authentication ldap ssl disable
Secure LDAP is disabled.
```

Configure LDAP server certificate verification with imported CA certificates

You can change the TLS request certificate behavior.

Steps

1. Change the TLS request certificate behavior by using the `authentication ldap ssl set tls_reqcert` command.

   Do not verify the certificate:

```
# authentication ldap ssl set tls_reqcert never
"tls_reqcert" set to "never". LDAP server certificate will not be verified.
```

   NOTE: If LDAP is configured for Active Directory, the TLS request certificate behavior cannot be set to `never`.

   Verify the certificate:

```
# authentication ldap ssl set tls_reqcert demand
"tls_reqcert" set to "demand". LDAP server certificate will be verified.
```

2. Reset the TLS request certificate behavior by using the `authentication ldap ssl reset tls_reqcert` command. The default behavior is `demand`:

```
# authentication ldap ssl reset tls_reqcert
tls_reqcert has been set to "demand". LDAP Server certificate will be verified with
imported CA certificate.Use "adminaccess" CLI to import the CA certificate.
```

Manage CA certificates for LDAP

You can import or delete certificates and show current certificate information.

Steps

1. Import a CA certificate for LDAP server certificate verification by using the `adminaccess certificate import` command.

   Specify `ldap` for `ca application`:

```
# adminaccess certificate import {host application {all | aws-federal | ddboost | https
| keysecure | dsm | ciphertrust | gklm | <application-list>} | ca application {all |
cloud | ddboost | ldap | login-auth | keysecure | dsm | rsa-securid | ciphertrust | gklm
| <application-list>}} [file <file-name>]
```

2. Delete a CA certificate for LDAP server certificate verification by using the `adminaccess certificate delete` command.

   Specify `ldap` for `application`:

```
# adminaccess certificate delete {subject <subject-name> | fingerprint <fingerprint>}
[application {all | aws-federal | cloud | ddboost | ldap | login-auth | https |
keysecure | dsm | ciphertrust | gklm | support | <application-list>}]
```

3. Show current CA certificate information for LDAP server certificate verification by using the `adminaccess certificate show` command:

```
# adminaccess certificate show imported-ca application ldap
```

Viewing NIS authentication information

The NIS Authentication panel displays the NIS configuration parameters and whether NIS authentication is enabled or disabled.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the NIS Authentication panel.

Enabling and disabling NIS authentication

Use the NIS Authentication panel to enable and disable NIS authentication.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the NIS Authentication panel.
3. Click **Enable** next to NIS Status to enable or **Disable** to disable NIS Authentication.

   The Enable or Disable NIS dialog box appears.

4. Click **OK**.

Configuring the NIS domain name

Use the NIS Authentication panel to configure the NIS domain name.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the NIS Authentication panel.
3. Click **Edit** next to Domain Name to edit the NIS domain name.

   The Configure NIS Domain Name dialog box appears.

4. Enter the domain name in the **Domain Name** box.
5. Click **OK**.

Specifying NIS authentication servers

Use the NIS Authentication panel to specify NIS authentication servers.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the NIS Authentication panel.
3. Below Domain Name, select one of the following:
   - **Obtain NIS Servers from DHCP** The system automatically obtains NIS servers using DHCP
   - **Manually Configure** Use the following procedures to manually configure NIS servers.
   - To add an authentication server, click Add (**+**) in the server table, enter the server name, and click **OK**.

- To modify an authentication server, select the authentication server name and click the edit icon (pencil). Change the server name, and click **OK**.
- To remove an authentication server name, select a server, click the X icon, and click **OK**.

4. Click **OK**.

Configuring NIS groups

Use the NIS Authentication panel to configure NIS groups.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the NIS Authentication panel.
3. Configure the NIS groups in the NIS Group table.
   - To add a NIS group, click Add (**+**), enter the NIS group name and role, and click **Validate**. Click **OK** to exit the add NIS group dialog box. Click **OK** again to exit the **Configure Allowed NIS Groups** dialog box.
   - To modify an NIS group, select the checkbox of the NIS group name in the NIS group list and click Edit (pencil). Change the NIS group name, and click **OK**.
   - To remove an NIS group name, select the NIS group in the list and click Delete **X**.
4. Click **OK**.

Configuring SSO authentication

The Single Sign-On (SSO) panel displays the SSO configuration parameters and whether SSO is enabled or disabled. Configuring SSO requires action on both the protection system and the SSO provider. SSO is supported on physical protection systems, and locally installed DD VE instances. Cloud-based DD VE instances are not supported.

About this task

SSO allows you to register a protection system with a supported SSO provider to use the SSO provider credentials for system-level user authentication. Logging in using single sign-on (SSO) describes how to log in using SSO after SSO is configured, an SSO user group is created, and SSO is enabled.

NOTE: Data Protection Central (DPC) is the only supported SSO provider. DPC version 19.1 is required to use SSO.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Single Sign-On (SSO) panel.

Registering the protection system in Data Protection Central (DPC)

About this task

Complete the following steps to register the protection system in DPC.

Steps

1. Log in to the DPC and navigate to the **System Management**.
2. Add the system to DPC.

   NOTE: DPC requires sysadmin credentials for the system.

3. Refresh the Single Sign-On (SSO) panel in DD SM to confirm that the system is registered with DPC.

Enabling and disabling SSO

Use the Single Sign-On (SSO) panel to enable or disable SSO.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Single Sign-On (SSO) panel.
3. Click **Enable** next to Single Sign-On Status to enable or **Disable** to disable SSO.

   The Enable or Disable SSO dialog box appears.

4. Click **OK**.

Configuring Single Sign-On (SSO) groups

Use the Single Sign-On (SSO) panel to configure SSO user groups.

About this task

At least one SSO user group is required to use SSO functionality.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Single Sign-On (SSO) panel.
3. Configure the SSO user groups in the table.
   - To add an SSO user group, click Add (**+**), enter the SSO user group name and domain name, select the management role, and click**OK**.

     NOTE: Admin users can set a group management role to user, admin, backup-operator, or limited-admin. Limited-admin users can set a group management role to user or backup operator.

     NOTE: If a group name belongs to multiple domains, set up the same group name with all domain names on the protection system with the desired role, or make sure the domain name the user will log in with is configured on system with the desired role. This is important for Active Directory configurations with child or sub domains.
   - To modify an SSO user group, select the checkbox of the group name in the SSO group list and click Edit (pencil). Change the management role and click **OK**.
   - To remove an SSO user group, select the group in the list and click Delete (**X**).

Configuring and enabling multifactor authentication (MFA)

Multifactor authentication adds an extra layer of security on the protection system by requiring the security officer and system administrator to enter an RSA SecurID passcode before logging into the system (MFA for login), and certain destructive commands or configuration changes (MFA for security officer oversight) are allowed. The **Multifactor Authentication** panel provides the ability to configure, enable, edit, and disable multifactor authentication.

Prerequisites

Before configuring MFA, verify Telnet is disabled on the protection system. MFA and Telnet cannot be enabled at the same time.

Add the protection system users to RSA Authentication Manager. The white paper *PowerProtect Appliances: Configuring and Enabling Multifactor Authentication for Local Users*, available from the Dell Technologies Info Hub (https://infohub.delltechnologies.com/), provides additional information about configuring RSA Authentication Manager for PowerProtect DD systems.

The following DD-specific requirements apply:

- For local, NIS, or AD users, add the user to the RSA internal database first.
- For LDAP users:
  - Add the external identify source to the RSA Operation Console.
  - Link the external identity source in the RSA Security Console.
- Create a unique user ID for each user:
  - For local users, create the user ID in the format *<User-ID>@<DD-serial-number>*.

    NOTE: Run the `system show serialno command` to get the system serial number.
  - Do not append the system serial number to the user IDs for AD or NIS users.

About this task

MFA for login is only supported for username and password login from the GUI or SSH. it is not supported for certificate, or token-based login. RSA SecurID is the only supported MFA server.

To ensure that backup applications can access the system without a passcode, MFA for login provides an option to disable MFA for the sysadmin user only.

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.
2. Expand the Multi-factor Authentication panel.
3. Click **Configure**.

   The RSA SecurID Server Authentication dialog box appears.
4. Specify the RSA configuration values:
   a. In the **Server URL** field, specify the RSA server URL.
   b. In the **Client Key** field, specify the RSA client key.
   c. In the **Client ID** field, specify the RSA client ID.
   d. In the **Connection Timeout** field, optionally change the connection timeout value.
   e. In the **Read Timeout** field, optionally change the Read timeout value.
   f. In the **Replica URL(s)** field, optionally specify any replica URLs for the RSA server.
5. Click **OK**.
6. Click **+ Add** to select and add the RSA server certificate to the protection system.
7. Click **OK**.
8. Click **Enable**.
9. Specify the security officer credentials:
   a. In the **Username** field, specify the security officer username.
   b. In the **Password** field, specify the security officer password.
10. Click **Next**.
11. In the **Password** field, specify the sysadmin password.
12. Click **Finish**.
13. Test the connection to the RSA SecurID server: Click **Test Connection**.

    NOTE: Testing the connection is mandatory for MFA for login. If the connection is not tested, sysadmin and security officer users cannot log in to the system. Testing the connection for other users is recommended but not required.

    a. Click **Test Connection**.
    b. In the **Username** field, specify the username to test.
    c. In the **Passcode** field, specify the sysadmin RSA passcode.
    d. Click **OK**.
14. Dell recommends creating an MFA troubleshooting user for situations where the RSA passcode does not work, and access to the system is required to disable MFR or perform troubleshooting steps.
15. If necessary, click **Edit** to make changes to the RSA configuration values.

Disabling multifactor authentication

Steps

1. Select **Administration** > **Access** > **Authentication**.

   The Authentication view appears.

2. Expand the Multifactor Authentication panel.
3. Click **Disable**.

   The Disable RSA SecurID dialog box appears.

4. Specify the security officer credentials:
   a. In the **Username** field, specify the security officer username.
   b. In the **Password** field, specify the security officer password.
   c. In the **Passcode** field, specify the security officer RSA passcode.
5. Click **Next**.
6. Specify the sysadmin credentials:
   a. In the **Password** field, specify the sysadmin password.
   b. In the **Passcode** field, specify the sysadmin RSA passcode.
7. Click **Finish**.

Diagnosing authentication issues

DDOS provides the ability to diagnose authentication issues for Active Directory from within the DD System Manager interface.

Steps

1. Select **Administration** > **Access** > **Authentication**
2. Expand the Active Directory/Kerberos Authentication panel.
3. Click **Diagnose**.
4. Select an issue to investigate, and click **Diagnose**.
5. Provide the requested information.

   To diagnose issues logging in as an Active Directory user, provide:
   - Active Directory server IP address
   - Active Directory server FQDN
   - Active Directory service username

     NOTE: The Active Directory user account specified here requires the following privileges:
     - Read-only access to the base DN identified by the domain name.
     - Read-only access to query attributes of all users in the base DN.
     - Read-only access to query attributes of the machine account for the protection system.
   - Active Directory service password
   - Username experiencing login failure

   To diagnose issues joining the system to an Active Directory Domain, provide:
   - Active Directory server IP address
   - Active Directory server FQDN
   - Active Directory service username
   - Active Directory service password
6. Click **Diagnose**.
7. View the report.
   - Click **View Report** to view the report online. Each item in the Action Items table can be clicked for additional details.
   - Click **Download** to download a copy of the report.
8. Review and implement the suggested fixes for the issue, and retry the operation.

Change system authentication method

The protection system supports password-based login, or certificate-based authentication. Password-based authentication is the default method.

Prerequisites

Certificate-based authentication requires SSH keys and CA certificates are imported to allow users to authenticate with the system when password-based login is disabled.

About this task

Complete the following steps to change the system authentication method from password-based login to certificate-based authentication.

Steps

1. Select **Administration** > **Access**.
   The Access Management view appears.
2. Click **Manage CA Certificates**.
3. Click **Add** to create a new certificate.
4. Add the certificate.
   - Select **I want to upload the certificate as a .pem file** and click **Choose File** to select the certificate file and upload it to the system.
   - Select **I want to copy and paste the certificate text** to copy and paste the certificate text into the text field.
5. Click **Add**.
6. Select **More Tasks** > **Change Login Options**.
7. In the **Password Based Login** drop-down menu, select **Disable**.

   NOTE: The drop-down menu is disabled if the required SSH keys and CA certificates are not configured on the system

8. Click **OK**.
   If a security policy is configured, the system prompts for security officer credentials. Provide the credentials and click **OK**.

Reset the system authentication method to password-based authentication.

About this task

Complete the following steps to change the system authentication method from certificate-based authentication to password-based authentication.

Steps

1. Select **Administration** > **Access**.
   The Access Management view appears.
2. Select **More Tasks** > **Change Login Options**.
3. In the **Password Based Login** drop-down menu, select **Enable**.
4. Click **OK**.
   If a security policy is configured, the system prompts for security officer credentials. Provide the credentials and click **OK**.

Reset the iDRAC password

If the iDRAC password for DD3300, DD6400, DD6900, DD9400, and DD9900 systems is lost or forgotten, it is possible to reset the password to the factory default setting.

About this task

The Data Domain system requires that the Integrated Dell Remote Access Controller (iDRAC) is configured for system upgrade and maintenance operations. Additionally, the system supports the use of iDRAC to change security settings, and remotely power the system on and off.

CAUTION: Do not use iDRAC to change the storage configuration, system settings, or BIOS settings, as making changes will impact system functionality. Contact Support if changes are required in any of these areas.

Steps

1. Connect to the system serial console or connect KVM to the system.
2. Reboot the system.
3. During the system boot process, press **F2** to access the BIOS menu.
4. Select **iDRAC Settings**.
5. Select **Reset iDRAC configurations to defaults all**.
6. Select **Yes** to confirm the reset.
7. Select **Continue**.
8. Exit the BIOS and reboot.

Results

The iDRAC configuration resets to the following username and password:

- Username: root
- Password: calvin

Viewing active users

The Active Users tab displays the names of users who are logged into the system and statistics about the current user sessions.

Steps

Select **Administration** > **Access** > **Active Users**.

The Active Users list appears and displays information for each user.

NOTE: To manage local users, click **Go to Local Users**.

Related concepts

Local user account management

9

Upgrading the System

This chapter presents the following topics:

Topics:

Managing system upgrades

To ensure that the state of your system can accommodate a successful upgrade, complete the preparation steps outlined in the following sections. After you have verified that the upgrade can complete, transfer the software to the system to be upgraded, and then start the upgrade.

The upgrade can take 90 minutes or more based on the DDOS version, system model, and storage attached.

HA system upgrades

The upgrade process on an HA system automatically upgrades both the active and standby nodes.

The HA nodes are upgraded one at a time. The standby node is rebooted and upgraded first. The newly upgraded standby node then takes over the active role through an HA failover. After the failover, the second node is rebooted and assumes the role of the standby node after the upgrade.

For an HA system, transfer the software to the active node and start the upgrade from the active node. Use the floating IP address to access DD System Manager to perform software upgrades.

System upgrade operations that require data conversion cannot start until both systems are upgraded to the same level and HA state is fully restored.

Minimally disruptive upgrade

The minimally disruptive upgrade (MDU) feature lets you upgrade specific software components or apply bug fixes without needing to perform a system reboot.

The MDU feature can prevent significant downtime during software upgrades, because the upgrade disrupts only those services that depend on the component being upgraded.

Not all software components qualify for an MDU; some components must be upgraded as part of a regular DDOS system software upgrade.

- A DDOS software upgrade uses a large RPM (upgrade bundle), which performs upgrade actions for all of the components of DDOS.
- MDU uses smaller component bundles, which upgrade specific software components individually. The MDU RPM is much smaller than a full DDOS RPM.

Contact Support to determine if an MDU is available for a specific issue.

Support software

Support software is an MDU package created and signed by Support Engineering to address specific issues.

By default, the DD system does not allow support software to be installed on the system. Contact Support for more information about support software.

Viewing and obtaining upgrade packages

DD System Manager allows you to view and manage up to five upgrade packages on the DDOS system. To upgrade the system, download an upgrade package from the Online Support site to a local computer, and then upload it to the target system.

About this task

NOTE: You can use FTP, NFS, or SCP to copy an upgrade package to a system. You can manage an unlimited number of system upgrade packages through a network share linked to the `/ddvar/releases` directory on the protection system. FTP is disabled by default. To use NFS, export and mount `/ddvar` from an external host.

Steps

1. Select **Maintenance** > **System**. To view the MD5 and SHA256 checksums, select an upgrade package and click **View Checksum**.
2. To obtain an upgrade package, click the **Dell Online Support** link, click **Downloads**, and use the search function to locate the recommended package for your system. Save the upgrade package to the local computer.
3. If more than four packages are listed in Upgrade Packages Available, remove at least one package before uploading the new package.
4. Click **Upload Upgrade Package**.
5. In the **Upload Upgrade Package** dialog box, click **Browse**, navigate to and select the file, and click **Open**.
6. Click **OK**.
   The `.rpm` file is downloaded and appears in the **Upgrade Packages Available** list.
7. To verify the upgrade package integrity, click **View Checksum** and compare the calculated checksum displayed in the dialog box to the authoritative checksum on the Online Support site.
8. To manually initiate an upgrade precheck, select an upgrade package and click **Upgrade Precheck**.

Preparing the system for upgrade

Perform these manual tasks before attempting an upgrade to avoid potential upgrade failure. Complete the upgrade pre-checks described in this section approximately one week before the scheduled DDOS upgrade to allow time to address any issues found during the pre-checks.

Prerequisites

- Reboot the system.
- For HA systems, do not reboot the system until after completing the steps in this procedure.
- Review the section Troubleshooting upgrade errors for other issues that can cause failure.

About this task

These tasks are not performed automatically by any process.

Table 6. Prepare the system for upgrade

Step	Step description	Command
1.	Check for current alerts and address any disk or hardware failures before upgrading.	`# alert show current`
2.	Verify the system components are running the latest firmware versions.	`# enclosure show firmware`

Table 6. Prepare the system for upgrade (continued)

Step	Step description	Command
3.	Verify the system network configuration.	**# net show config** **# net show hardware**
4.	Ensure that all network interfaces are up and have appropriate IP addresses.	**# net show settings**
5.	Check the disk states, and do not perform the upgrade if the system is low on spares or has disks that show in the absent, failed, or reconstructing states.	**# disk show state**
6.	Check the disk reliability, and replace any disks that have more than 50 reallocated sectors.	**# disk show reliability-data**
7.	Check that the enclosure status is OK for all devices.	**# enclosure show all**
8.	Check whether the enclosure topology is correct. Check whether any error appears with an asterisk (*) next to the `enc.ctrl.port` field. Also check the **Error Message** field for any errors such as `A possible problem was detected for this shelf controller or the cable connected to it.`	**# enclosure show topology**
9.	Check that the device port mapping is correct.	**# system show hardware**
10.	Check the link speed for connected ports.	**# system show ports**
11.	Check the status of the file system to determine that file system is enabled and running normally.	**# filesys status**
12.	Check the replication status.	**# replication status**
13.	For systems with AWS cloud units upgrading to DDOS 7.4 and higher, the change in endpoint format requires an update to the network infrastructure to allow the system to remain connected to the endpoint after the upgrade is complete.	Change the endpoint format from `s3-<region>.amazonaws.com` to **s3.<region>.amazonaws.com**.
14.	For systems with AWS cloud units upgrading to DDOS 7.8 and higher, the us-east-1 region no longer supports the legacy endpoint s3.amazonaws.com. The us-east-1 region now requires the endpoint s3.us-east-1.amazonaws.com. Verify the firewall is open to reach the new endpoint before upgrading to DDOS 7.8.	
15.	Check if any backup and restore activity is in progress, and if so, stop it.	**# system show stats**
16.	Run an Autosupport Report just prior to performing the DDOS upgrade to determine if the system reports errors that need to be resolved before the upgrade.	**# autosupport send <your_email_address>**
17.	If the Autosupport Report indicated issues with the system, check `kern.info` log, and if you notice frequent failures in hardware, contact Support to inspect your system before you perform the upgrade. Search for the string `ERROR` in the log file.	**# log view debug/platform/kern.info**
18.	For systems with any X.509 version 4 CA certificates in its trust-store, the upgrade is blocked. Recreate the trust relationships with CA certificates using X.509 version 3.	• `adminaccess trust add` • `ha create` • `replication add` • `managed-system add`

Next steps

For HA systems, follow the reboot instructions described in the "Upgrade considerations for HA systems" section.

Automatic tasks performed by the upgrade script (in the .rpm file) prior to upgrade

These tests precede the actual system upgrade process. The system:

1. Determines whether two different kinds of NVRAM cards are present.
2. Checks the `/ddr` partition and `/` (root) partition sizes for space utilization.
3. Checks the OST version.
4. Determines whether the RAID metagroup is assembled. If it is not assembled, the upgrade process does not begin.
5. Determines available space for the file system.
6. Determines whether sufficient space is available for the upgrade.
7. Checks the VTL version, if VTL is present.
8. Determines whether the file system is enabled, and if it is not enabled, enable it.
9. Determines whether VTL is enabled.
10. Checks the VTL pools to ensure that they can be converted to MTrees.
11. Determines whether sufficient VTL space is available.
12. Ensures that the numbers of MTrees and VTL pools do not exceed 100.
13. Determines whether all dg0 disks are located on head unit. If not, the upgrade process does not begin, and the problem must be addressed.
14. Determines whether the file system can be shutdown without problems. If the file system cannot be shut down in a clean manner, the upgrade process will stop.
15. For DD6400, DD6900, DD9400, and DD9900 systems only: Determines whether any iDRAC administrator accounts are present on the system.
16. Checks if file system cleaning is running, and displays a warning if a cleaning operation is in progress.
17. Checks if cloud cleaning is running, and displays a warning if a cleaning operation is in progress.
18. Checks if data movement to the cloud is running, and displays a warning if a data movement operation is in progress.

The upgrade process quits if it encounters a failure in any of the tasks listed.

Upgrading the DD system using DD System Manager

When an upgrade package file is present on the system, you can use DD System Manager to perform an upgrade using that upgrade package.

Prerequisites

- Read the DDOS Release Notes for the complete upgrade instructions and coverage of all the issues that can impact the upgrade.
- Log out of any CLI sessions on the system where the upgrade is to be performed.
- Reboot the system to verify that the hardware is in a clean state. Resolve any issues discovered during the reboot. For an MDU upgrade, a reboot might not be needed.

About this task

NOTE: This topic assumes that you are updating only DDOS. If you make hardware changes, such as adding, swapping, or moving interface cards, you must update the DDOS configuration to correspond with the hardware changes.

When upgrading from DDOS 7.6 or earlier to DDOS 7.7 or later, the system will automatically roll back the upgrade to the previously installed working version if the system encounters an error that causes the upgrade to fail.

Steps

1. Log into DD System Manager on the protection system where the upgrade is to be performed.
2. Select **Maintenance** > **System**.
3. From the **Upgrade Packages Available** list, select the package to use for the upgrade.

   NOTE: You must select an upgrade package for a newer version of DDOS. DDOS does not support downgrades to previous versions.

4. Click **Perform System Upgrade**.

The **System Upgrade** dialog box appears and displays information about the upgrade and a list of users who are currently logged in to the system to be upgraded.

5. Verify the version of the upgrade package, and click **OK** to continue with the upgrade.

   The **System Upgrade** dialog box displays the upgrade status and the time remaining.

   Wait for the upgrade to complete before using DD System Manager to manage the system. If the system restarts, the upgrade might continue after the restart, and DD System Manager displays the upgrade status after login. If possible, keep the System Upgrade progress dialog box open until the upgrade completes or the system restarts. A Login link appears when the upgrade is complete.

   NOTE: To view the status of an upgrade using the CLI, enter the `system upgrade status` command. Log messages for the upgrade are stored in `/ddvar/log/debug/platform/upgrade-error.log` and `/ddvar/log/debug/platform/upgrade-info.log`.

6. If the system powers down, you must remove AC power from the system to clear the prior configuration. Unplug all power cables for 30 seconds, and then plug them back in. The system powers on and restarts.
7. If the system does not automatically power on and there is a power button on the front panel, press the button.

Next steps

The following requirements might apply after completing an upgrade.

- For environments that use self-signed SHA256 certificates, do the following:
  1. Run the **`adminaccess certificate generate self-signed-cert regenerate-ca`** command to regenerate the self-signed CA and host certificates. Regenerating the certificates breaks existing trust relationships with external systems.

     NOTE: The default validity of the self-signed host certificate is 12 months.

  2. Run the **`adminaccess trust add host hostname type mutual`** command to reestablish mutual trust between the protection system and the external system.
- If the system shows existing or configured FC ports with missing WWPN or WWNN information, or reports that no FC host bus adapter (HBA) driver is installed, run the `scsitarget endpoint enable all` command.
- For environments running MTree analytics, set the frequency if required, as the upgrade reset the Active Tler frequency to one week.
- If DD Boost encryption is set to none, upgrading to DDOS 7.12 or higher triggers a security alert to remind the user to enable on-the-wire encryption.

NOTE: With collection replication, no files are visible on the destination system if replication was not finished before starting the upgrade. After the upgrade, wait until replication completes to see files on the destination system.

Upgrading HA systems

HA systems require additional steps before and after the upgrade operation.

Prerequisites

Perform the manual checks described in Preparing the system for upgrade before rebooting the HA system.

If the customer system has MTrees where the `attribute_btree_layout` version is set to `OLD`, Dell recommends performing the upgrade in local mode.

About this task

When upgrading an HA system, upload the upgrade RPM package to the active node.

Steps

1. Run the `ha status` command to verify the HA system state.

   The HA system must be in a highly available state, with both nodes online before performing the DDOS upgrade.

```
# ha status
HA System Name: dd9900-ha3a.example.com
HA System Status: highly available
```

```
Node Name                    Node ID   Role      HA State
-------------------------- --------- --------- --------
dd9900-ha3a-p0.example.com   0       active    online
dd9900-ha3a-p1.example.com   1       standby   online
-------------------------- --------- --------- --------
```

2. Reboot the standby node (node 1).
3. Run the `ha status` command to verify that the HA system status is `highly available`.
4. Run the `ha failover precheck` command to verify that the HA system status is ready for failover.
5. Run the `ha failover` command to initiate a failover from the active node to the standby node.
6. Run the `ha status` command to verify that node 1 is the active node and node 0 is the standby node.

```
# ha status
HA System Name: dd9900-ha3a.example.com
HA System Status: highly available
Node Name                    Node ID   Role      HA State
-------------------------- --------- --------- --------
dd9900-ha3a-p0.example.com   0       standby   online
dd9900-ha3a-p1.example.com   1       active    online
-------------------------- --------- --------- --------
```

7. Reboot the standby node (node 0).
8. Run the `ha status` command to verify that the HA system status is `highly available`.
9. Run the `ha failover precheck` command to verify that the HA system status is ready for failover.
10. Run the `ha failover` command to initiate a failover from the active node to the standby node.
11. Run the `ha status` command to verify that node 0 is the active node and node 1 is the standby node.

```
# ha status
HA System Name: dd9900-ha3a.example.com
HA System Status: highly available
Node Name                    Node ID   Role      HA State
-------------------------- --------- --------- --------
dd9900-ha3a-p0.example.com   0       active    online
dd9900-ha3a-p1.example.com   1       standby   online
-------------------------- --------- --------- --------
```

12. Initiate the upgrade from the active node.
    DDOS automatically recognizes the HA system and performs the upgrade procedure on both nodes. The HA upgrade runs in the following sequence:
    a. The standby node is upgraded, then reboots.
    b. The HA system initiates a failover and the standby node takes over as the active node.
    c. The original active node is upgraded, then reboots and remains as the standby node.

Results

After both nodes are upgraded, the system does not perform another failover to return the nodes to their original configuration.

Next steps

After the upgrade procedure is complete, run the `ha status` command again to verify that the system is in a highly available state, and both nodes are online.

Optionally run the `ha failover` command to return the nodes to their pre-upgrade roles.

Removing an upgrade package

Use this procedure to make room for a new upgrade package.

Steps

1. Select **Maintenance** > **System**.

2. From the list titled **Upgrade Packages Available on *<protection system>***, select the package to remove. One package can be removed at a time.
3. Click **Remove Upgrade Package**.

Troubleshooting upgrade errors

DDOS upgrade performs certain tests before a system upgrade. If a test fails, the upgrade could fail.

The following circumstances could cause the upgrade to fail:

Running tasks

- A replication initialization is in progress.
- The file system did not shutdown cleanly, resulting in a core dump.
- A previous upgrade did not complete successfully.

Configuration issues

- The system is not configured correctly. For example, NFS mount points were manually created under root.
- Inspection of digests and signatures that are contained in `.rpm` file indicates that the signature is not valid. The valid signatures are SHA1 or MD5.
- NFS mount points are unknown.
- The RAID metagroup is not assembled.
- All dg0 disks are not located on the head unit.
- The file system cannot be shut down in a clean manner.
- Two different kinds of NVRAM cards are present.
- The OST version is not compatible.
- A VTL version is incompatible.
- The file system is not enabled and cannot be enabled automatically.
- VTL is not enabled.
- A check of the VTL pools indicates that they cannot be converted to MTrees.
- Features are incompatible between the original and target versions of DDOS:
  - DDOS versions 7.0 and later do not support DD Extended Retention. Any system with DD Extended Retention enabled cannot be upgraded to DDOS 7.0. Migrate the data to a system running DDOS 6.1.X or 6.2.X without DD Extended Retention running. After the migration, upgrade the DDOS 6.1.X or 6.2.X system without DD Extended Retention to DDOS 7.0.
  - DDOS versions 7.0 and later do not support the RSA key manager for data at rest encryption. If a DDOS 6.X system uses the RSA key manager for data at rest encryption, upgrades to DDOS 7.0 and later are not permitted. Disable data at rest encryption to proceed with the DDOS upgrade.
  - DDOS versions 7.0 and later use a newer version of OpenSSH than DDOS versions 6.2 and earlier. Upgrades to DDOS 7.0 will overwrite unsupported cipher-list and mac-list values with the default values for DDOS versions 7.0 and later, and the commands to set cipher-list and mac-list values will fail with error messages when unsupported values are specified after the upgrade.

Space issues

- Storage is functionally deficient, such as an enclosure is missing.
- The `/` (root) or `/ddr` partition is full with log files, core dumps, and so forth.
- Available space for the file system is insufficient.
- Available space for the upgrade is insufficient.
- Available space for VTL is insufficient.
- The number of MTrees or VTL pools exceeds 100.

Slow connection

NOTE: For security reasons, there is a 30-minute time limit for the upload of RPM packages for DDMC and DD system upgrades using the DDMC GUI. If you have a slow connection from a client machine to the DDMC and the upload takes more than 30 minutes, the connection drops and you cannot use DDMC to upload the package.

Workaround: Use the CLI to upload the package into DDMC (for example, use `SCP/PSCP` from a Unix terminal or Windows CMD).

For DDMC upgrades, upload the package to `/ddr/var/releases`.

For DD System upgrades, upload the package to `/ddr/var/ddr-releases`.

10

Managing Network Connections

This chapter presents the following topics:

Topics:

Managing network connections

Network connection management features allow you view and configure network interfaces, general network settings, and network routes.

Managing HA system network connections

The HA system relies on two different types of IP addresses: fixed and floating. Each type has specific behaviors and limitations.

On an HA system, fixed IP addresses:

- Are used for node management via the CLI
- Are attached ("fixed") to the node
- Can be static or DHCP, IPv6 SLAAC, and IPv6 Linklocal

  NOTE: SLAAC and Linklocal are auto-generated address that appear when the interface is in a running state. Users have no control over these addresses, but they are available to transfer data.
- Configuration is done on the specific node with the `type fixed` argument

NOTE: All file system access should be done through a floating IP address.

Floating IP addresses exist only in the two-node HA system. During failover, the IP addresses "float" to the new active node and are:

- Only configured on the active node
- Used for file system access and most configuration
- Can only be static
- Configuration requires the `type floating` argument

Managing network interfaces

You can manage the physical interfaces that connect the system to a network and create logical interfaces to support link aggregation, load balancing, and link or node failover.

Viewing interface information

The Interfaces tab enables you to manage physical and bonded interfaces, VLANs, DHCP, DDNS, and IP addresses and aliases.

About this task

Consider the following guidelines when managing IPv6 interfaces.

- MTree replication are supported over IPv6 networks, which allows you to take advantage of the IPv6 address space. Simultaneous replication over IPv6 and IPv4 networks is also supported, as is Managed File Replication using DD Boost.
- There are some restrictions for interfaces with IPv6 addresses. For example, the minimum MTU is 1280. If you try to set the MTU lower than 1280 on an interface with an IPv6 address, an error message appears and the new MTU size is rejected. If you try to set an IPv6 address on an interface with an MTU lower than 1280, an error message appears. An IPv6 address set on a VLAN which is on a physical or bonded interface will impact that physical or bonded interface. If the MTU of the physical or bonded interface is given an MTU less than 1280 but it has a VLAN with an IPv6 address on it, the setting of the MTU will be rejected with an error message. In addition, if the physical or bonded interface has an MTU less than 1280 and an associated VLAN interface is given an IPv6 address, that address is rejected because the base MTU is too small.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. To filter the interface list by interface name, enter a value in the **Interface Name** field and click **Update**.

   Filters support wildcards, such as eth*, veth*, or eth1*
3. To filter the interface list by interface type, select a value from the Interface **Type** menu and click **Update**.

   On an HA system, there is a filter dropdown to filter by IP Address Type (Fixed, Floating, or Interconnect).
4. To return the interfaces table to the default listing, click **Reset**.
5. Select an interface in the table to populate the Interface Details area.
6. To view IPMI interface configuration and management options, click **View IPMI Interfaces**.

   This link displays the **Maintenance** > **IPMI** information.

Physical interface names

The layout of physical interface names varies on different protection systems and option cards.

The physical interface name format is `ethXy`, where *X* is the slot number for an on-board port or an option card and *y* is an alphanumeric string. For example, eth0a.

- DD6300, DD6800, and DD9300 systems provide one on-board 1G Base-T NIC port: ethMa.
- DD9800 systems provide four on-board 1G Base-T NIC ports: ethMa (bottom left), ethMb (top left), ethMc (bottom right), and ethMd (top right).
- For most vertical I/O module NIC interfaces, the port numbering goes from top to bottom, with eth*X*a at the top.
- For most horizontal I/O module NIC interfaces, the port numbering goes from left to right, with eth*X*a on the left.
- The horizontal I/O module slots on the left-hand side (expansion riser 1) of the DD3300, DD6400, DD6900, DD9400, and DD9900 systems are inverted. The port numbering on the I/O modules in these slots goes from right to left, with eth*X*a on the right.

General interface configuration guidelines

Review the general interface configuration guidelines before configuring system interfaces.

- When supporting both backup and replication traffic, if possible, use different interfaces for each traffic type so that neither traffic type impacts the other.
- When replication traffic is expected to be less than 1 Gb/s, if possible, do not use 10 GbE interfaces for replication traffic because 10 GbE interfaces are optimized for faster traffic.
- If a service uses a non-standard port and the user wants to upgrade to DDOS 7.X, or the user wants to change a service to use a non-standard port on a DDOS 7.X system, add a net filter function for all the clients using that service to allow the client IP addresses to use the new port.

  NOTE: This happens automatically if the net filter auto function is active. The system detects the port is in use and adds a rule for it.
- For systems that use IPMI, if possible, reserve interface ethMa for IPMI traffic and system management traffic (using protocols such as HTTP, Telnet, and SSH). Backup data traffic should be directed to other interfaces.

  NOTE: This does not apply to DD3300, DD6400, DD6900, DD9400, and DD9900 systems which incorporate IPMI functions into iDRAC.

Configuring physical interfaces

Configure the management interface during initial system setup before configuring interfaces for user traffic. You must configure at least one physical interface before the system can connect to a network.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. Select an interface to configure.
3. Click **Configure**.
4. In the Configure Interface dialog box, determine how the interface IP address is to be set:

   NOTE: On an HA system, the Configure Interface dialog box has a field for whether or not to designate the Floating IP (Yes/No). Selecting **Yes** the `Manually Configure IP Address` radio button is auto-selected; Floating IP interfaces can only be manually configured.

   - Use DHCP to assign the IP address—in the IP Settings area, select **Obtain IP Address using DHCP** and select either **DHCPv4** for IPv4 access or **DHCPv6** for IPv6 access.

     Setting a physical interface to use DHCP automatically enables the interface.

     NOTE: If you choose to obtain the network settings through DHCP, you can manually configure the hostname at **Hardware** > **Ethernet** > **Settings** or with the `net set hostname` command. You must manually configure the host name when using DHCP over IPv6.
   - Specify IP Settings manually—in the IP Settings area, select **Manually configure IP Address**.

     The **IP Address** and **Netmask** fields become active.

5. If you chose to manually enter the IP address, enter an IPv4 or IPv6 address. If you entered an IPv4 address, enter a netmask address.

   NOTE: You can assign just one IP address to an interface with this procedure. If you assign another IP address, the new address replaces the old address. To attach an additional IP address to an interface, create an IP alias.

6. Specify Speed/Duplex settings.

   NOTE: Speed and duplex cannot be set on DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

   The combination of speed and duplex settings define the rate of data transfer through the interface. Select one of these options:

   - **Autonegotiate Speed/Duplex** — Select this option to allow the network interface card to autonegotiate the line speed and duplex setting for an interface. Autonegotiation is *not* supported on the following DD6300, DD6800, DD9300, and DD9800 I/O modules:
     - Dual port 10GbE SR Optical with LC connectors (using SFPs)
     - Dual port 10GbE Direct Attach Copper (SFP+ cables)
     - Quad port 2 port 1GbE Copper (RJ45) /2 port 1GbE SR Optical
     - **Autonegotiate Speed/Duplex** is required for all I/O modules on the DD3300, DD6400, DD6900, DD9400, and DD9900 systems:
       - Quad port 10GbE Base-T
       - Quad port 10GbE SFP+
       - Dual port 10/25GbE SFP28
       - Dual port 100GbE QSFP28
   - **Manually configure Speed/Duplex** — Select this option to manually set an interface data transfer rate. Select the speed and duplex from the menus.

     NOTE: This option is not available on DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

     - Duplex options are half-duplex, full-duplex, and unknown.
     - Speed options listed are limited to the capabilities of the hardware device. Options are 10 Mb, 100 Mb, 1000 Mb (1 Gb), 10 Gb, and unknown. The 10G Base-T hardware supports only the 100 Mb, 1000 Mb and 10 Gb settings.
     - Half-duplex is only available for 10 Mb and 100 Mb speeds.
     - 1 Gb and 10 Gb line speeds require full-duplex.

- The default setting for 10G Base-T interfaces is Autonegotiate Speed/Duplex. If you manually set the speed to 1000 Mb or 10 Gb, you must set the Duplex setting to Full.

7. Specify the MTU (Maximum Transfer Unit) size for the physical (Ethernet) interface.

   Do the following:

   - Click the **Default** button to return the setting to the default value.
   - Ensure that all of your network components support the size set with this option.

8. Optionally, select **Dynamic DNS Registration**.

   Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

   The DDNS must be registered to enable this option.

   NOTE: This option disables DHCP for this interface.

9. Click **Next**.

   The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state, which are applied after you click Finish.

10. Click **Finish** and **OK**.

MTU size values

The MTU size must be set properly to optimize the performance of a network connection. An incorrect MTU size can negatively affect interface performance.

Supported values for setting the maximum Transfer Unit (MTU) size for the physical (Ethernet) interface range from 600 to 9000 for IPv4, and 1280 to 9000 for IPv6. For 100 Base-T and gigabit networks, 1500 is the standard default.

NOTE: The minimum MTU for IPv6 interfaces is 1280. The interface fails if you try to set the MTU lower than 1280.

Moving a static IP address

A specific static IP address must be assigned to only one interface on a system. A static IP address must be properly removed from one interface before it is configured on another interface.

Steps

1. If the interface that hosts the static IP address is part of a DD Boost interface group, remove the IP address from that group.

   NOTE: Add the new IP address back to the DD Boost interface group after this task is complete.

2. Select **Hardware** > **Ethernet** > **Interfaces**.
3. Remove the static IP address that you want to move.
   a. Select the interface that is currently using the IP address you want to move.
   b. In the Enabled column, select **No** to disable the interface.
   c. Click **Configure**.
   d. Set the IP Address to 0.

      NOTE: Set the IP address to 0 when there is no other IP address to assign to the interface. The same IP address must not be assigned to multiple interfaces.

   e. Click **Next**, and click **Finish.**
4. Add the removed static IP address to another interface.
   a. Select the interface to which you want to move the IP address.
   b. In the Enabled column, select **No** to disable the interface.
   c. Click **Configure**.
   d. Set the IP Address to the match the static IP address you removed.

e. Click **Next**, and click **Finish**.
f. In the Enabled column, select **Yes** to enable the updated interface.

Bonded interface configuration guidelines

Bonded interface configuration guidelines apply to failover and aggregate bonded interfaces. There are additional guidelines that apply to either failover or aggregate interfaces but not both.

- The *virtual-name* must be in the form `vethx` where *x* is a number. The recommended maximum number is 99 because of name size limitations.
- You can create as many bonded interfaces as there are physical interfaces.
- Each interface used in a bonded interface must first be disabled. An interface that is part of a bonded interface is seen as disabled for other network configuration options.
- After a bonded interface is destroyed, the physical interfaces associated with it remain disabled. You must manually re-enable the physical interfaces.
- The number and type of cards installed determines the number of Ethernet ports available.
- Each physical interface can belong to one bonded interface.
- Bonded interfaces must be created from physical interfaces of the same. For example, all copper, all optical, all 1 Gb, or all 10 Gb. However, 1 Gb interfaces support bonding a mix of copper and optical interfaces. This applies to bonded interfaces across different cards with identical physical interfaces. The system checks for port speed mismatches when bonded interface is created, or a new physical interface is added to an existing bonded interface.

  NOTE: In DDOS 7.12 and higher, the system no longer checks the NIC vendor when adding interfaces from multiple NIC cards to a bonded interface.
- A system can support multiple mixed failover and aggregation bonded interfaces, subject to the restrictions above.
- Failover links improve network resiliency.
- Aggregate links improve network performance and resiliency by using two or more network interfaces in parallel, thus increasing the link speed for aggregated links and reliability over that of a single interface.
- Remove functionality is available using the **Configure** button. Click a bonded interface in the list of interfaces on the Interfaces tab and click **Configure**. From the list of interfaces in the dialog box, clear the checkbox for the interface to remove it from bonding (failover or aggregate), and click **Next**.
- For a bonded interface, the bonded interface is created with remaining secondary interfaces if the hardware for a secondary interface interface fails. If all the secondary interface fail, the bonded interface can no longer send and receive network traffic. This secondary interface hardware failure will generate managed alerts, one per failed secondary interface.

  NOTE: The alert for a failed secondary interface disappears after the failed secondary interface is removed from the system. If new hardware is installed, the alerts disappear and the bonded interface uses the new secondary interface after the reboot.
- On DD3300 systems, the ethMa interface does not support failover or link aggregation.

Guidelines for configuring a bonded interface for link aggregation

Link aggregation provides improved network performance and resiliency by using one or more network interfaces in parallel, thus increasing the link speed and reliability over that of a single interface. These guidelines are provided to help you optimize your use of link aggregation.

- It is recommended that you make interface changes only during scheduled maintenance downtime even though routing rules and gateways are reapplied after interface changes. Verify the routing configuration is still correct after making interface changes
- Enable aggregation on an existing bonded interface by specifying the physical interfaces and mode.
- Bring up the bonded interface and make sure an IP address is on the interface or an associated interface. The bonded interface may have VLANs and or aliases on it, each with an IP address, and therefore does not need an IP address directly on it to be fully functional.
- 1 GbE and 10 GbE interfaces cannot be aggregated together.
- 1 GbE copper and optical interfaces can be aggregated together.
- 10 GbE copper and optical interfaces cannot be aggregated together.

Guidelines for configuring a bonded interface for failover

Link failover provides improved network stability by identifying backup interfaces that can support network traffic when the primary interface is not operating. These guidelines are provided to help you optimize your use of link failover.

- A primary interface can be specified as part of the failover. If a primary interface removal is attempted from a failover, an error message appears.
- When the main secondary interface is used in a failover configuration, it must be explicitly specified and must also be a bonded interface to the bonded interface. If the primary interface goes down and multiple interfaces are still available, the next interface is randomly selected.
- All interfaces in a bonded interface must be on the same physical network. Network switches used by a bonded interface must be on the same physical network.
- The recommended number of physical interfaces for failover is greater than one. You can, however, configure one primary interface and one or more failover interfaces.

Bonded interface creation

Create a bonded interface to support link aggregation or failover. The bonded interface serves as a container for the links to be aggregated or associated for failover.

Creating a bonded interface for link aggregation

Create a bonded interface for link aggregation to serve as a container to associate the links that participate in aggregation.

About this task

A link aggregation interface must specify a link bonding mode and may require a hash selection. For example, you might enable link aggregation on bonded interface *veth1* to physical interfaces *eth2b* and *eth5c* in mode LACP (Link Aggregation Control Protocol) and hash XOR-L2L3.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the Interfaces table, disable the physical interface where the bonded interface is to be added by clicking **No** in the **Enabled** column.
3. From the **Create** menu, select **Virtual Interface**.
4. In the Create Virtual Interface dialog box, specify a bonded interface name in the **veth** box.

   Enter a bonded interface name in the form *vethx*, where *x* is a unique ID (typically one or two digits). A typical full bonded interface name with VLAN and IP Alias is *veth56.3999:199*. The maximum length of the full name is 15 characters. Special characters are not allowed. Numbers must be between 0 and 4094, inclusively.

5. In the **Bonding Type** list, select **Aggregate**.

   NOTE: Registry settings can be different from the bonding configuration. When interfaces are added to the bonded interface, the information is not sent to the bonding module until the bonded interface is brought up. Until that time the registry and the bonding driver configuration are different.

6. In the **Mode** list, select a bonding mode.

   Specify the mode that is compatible with the requirements of the system to which the interfaces are directly attached.

   - Round-robin

     Transmit packets in sequential order from the first available link through the last in the aggregated group.

   - Balanced

     Data is sent over interfaces as determined by the hash method selected. This requires the associated interfaces on the switch to be grouped into an Etherchannel (trunk) and given a hash via the Load Balance parameter.

   - LACP

     Link Aggregation Control Protocol is similar to Balanced, except that it uses a control protocol that communicates to the other end and coordinates which links within the bond are available for use. LACP provides a kind of heartbeat failover

and must be configured at both ends of the link. The LACP configuration can also specify whether the polling is fast (every second) or slow (every 30 seconds).

7. If you selected Balanced or LACP mode, specify a bonding hash type in the **Hash** list.

   Options are: XOR-L2, XOR-L2L3, or XOR-L3L4.

   XOR-L2 transmits through a bonded interface with an XOR hash of Layer 2 (inbound and outbound MAC addresses).

   XOR-L2L3 transmits through a bonded interface with an XOR hash of Layer 2 (inbound and outbound MAC addresses) and Layer 3 (inbound and outbound IP addresses).

   XOR-L3L4 transmits through a bonded interface with an XOR hash of Layer 3 (inbound and outbound IP addresses) and Layer 4 (inbound and outbound ports).

8. To select an interface to add to the aggregate configuration, select the checkbox that corresponds to the interface, and then click **Next**.

   The Create bonded interface *veth_name* dialog box appears.

9. Enter an IP address, or enter **0** to specify no IP address.
10. Enter a netmask address or prefix.
11. Skip the Speed/Duplex options as they are ignored for bonding.
12. Specify the MTU setting.
    - To select the default value (1500), click **Default**.
    - To select a different setting, enter the setting in the **MTU** box. Ensure that all of your network components support the size set with this option.
13. Optionally, select Dynamic DNS Registration option.

    Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

    The DDNS must be registered to enable this option.

14. Click **Next.**

    The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.

15. Click **Finish** and **OK**.

Creating a bonded interface for link failover

Create a bonded interface for link failover to serve as a container to associate the links that will participate in failover.

About this task

The failover-enabled bonded interface represents a group of secondary interfaces, one of which can be specified as the primary. The system makes the primary interface the active interface whenever the primary interface is operational. A configurable Down Delay failover option allows you to configure a failover delay in 900 millisecond intervals. The failover delay guards against multiple failovers when a network is unstable.

NOTE: Modifying up and down delay is not supported in DD3300. Only one default setting of 100ms is supported.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the interfaces table, disable the physical interface to which the bonded interface is to be added by clicking **No** in the **Enabled** column.
3. From the **Create** menu, select **Virtual Interface**.
4. In the Create Virtual Interface dialog box, specify a bonded interface name in the **veth** box.

   Enter a bonded interface name in the form `vethx`, where *x* is a unique ID (typically one or two digits). A typical full bonded interface name with VLAN and IP Alias is *veth56.3999:199*. The maximum length of the full name is 15 characters. Special characters are not allowed. Numbers must be between 0 and 4094, inclusively.

5. In the **Bonding Type** list, select **Failover**.

6. Select an interface to add to the failover configuration, and click **Next**. Virtual aggregate interfaces can be used for failover.

   The Create Virtual interface *veth_name* dialog box appears.

7. Enter an IP address, or enter **0** to specify no IP address.
8. Enter a netmask or prefix if an IP address was specified.
9. Skip the Speed/Duplex options as they are ignored for bonding.
10. Specify MTU setting.
    - To select the default value (1500), click **Default**.
    - To select a different setting, enter the setting in the MTU box. Ensure that all of your network path components support the size set with this option.
11. Optionally, select Dynamic DNS Registration option.

    Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

    The DDNS must be registered to enable this option.

    NOTE: This option disables DHCP for this interface.

12. Click **Next**.

    The Configure Interface Settings summary page appears. The values listed reflect the new system and interface state.

13. Complete the Interface, click **Finish** and **OK**.

Modifying a bonded interface

After you create a bonded interface, you can update the settings to respond to network changes or resolve issues.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the Interfaces column, select the interface and disable the bonded interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
3. In the **Interfaces** column, select the interface and click **Configure.**
4. In the **Configure Virtual Interface** dialog box, change the settings.
5. Click **Next** and **Finish.**

Configuring a VLAN

Create a new VLAN interface from either a physical interface or a bonded interface.

About this task

The recommended total VLAN count is 80. You can create up to 100 interfaces (minus the number of aliases, physical and bonded interfaces) before the system prevents you from creating any more.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the interfaces table, select the interface to which you want to add the VLAN.
3. Click **Create** and select**VLAN**.
4. In the Create VLAN dialog box, specify a VLAN ID by entering a number in the **VLAN ID** box. This is the tag the VLAN will use on the network.

   The range of a VLAN ID is between 1 and 4094 inclusive.

5. Enter an IP address, or enter **0** to specify no IP address.

   The Internet Protocol (IP) address is the numerical label assigned to the interface. For example, 192.168.10.23.

6. Enter a netmask or prefix.
7. Specify the MTU setting.

   The VLAN MTU must be less than or equal to the MTU defined for the physical or bonded interface to which it is assigned. If the MTU defined for the supporting physical or bonded interface is reduced below the configured VLAN value, the VLAN value is automatically reduced to match the supporting interface. If the MTU value for the supporting interface is increased above the configured VLAN value, the VLAN value is unchanged.

   - To select the default value (1500), click **Default**.
   - To select a different setting, enter the setting in the MTU box. DD System Manager does not accept an MTU size that is larger than that defined for the physical or bonded interface to which the VLAN is assigned.
8. Specify Dynamic DNS Registration option.

   Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

   The DDNS must be registered to enable this option.
9. Click **Next.**

   The **Create VLAN** summary page appears.
10. Review the configuration settings, click **Finish**, and click **OK.**

Modifying a VLAN interface

After you create a VLAN interface, you can update the settings to respond to network changes or resolve issues.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the **Interfaces** column, select the checkbox of the interface and disable the VLAN interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
3. In the Interfaces column, select the checkbox of the interface and click **Configure**.
4. In the **Configure VLAN Interface** dialog box, change the settings.
5. Click **Next** and **Finish**.

Configuring an IP alias

An IP alias assigns an additional IP address to a physical interface, a bonded interface, or a VLAN.

About this task

The recommended total number of IP aliases, VLAN, physical, and bonded interfaces that can exist on the system is 80. Although up to 100 interfaces are supported, as the maximum number is approached, you might notice slowness in the display.

NOTE: When using an HA pair, aliases cannot be created on the standby node. Create the alias on the active node then configure it on the standby node.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. Click **Create**, and select **IP Alias**.

   The Create IP Alias dialog box appears.
3. Specify an IP alias ID by entering a number in the **IP ALIAS Id** box.

   The range is 1 to 4094 inclusive.
4. Enter an IPv4 or IPv6 address.
5. If you entered an IPv4 address, enter a netmask address.
6. Specify Dynamic DNS Registration option.

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server. In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command.

The DDNS must be registered to enable this option.

7. Click **Next**.

   The Create IP Alias summary page appears.

8. Review the configuration settings, click **Finish**, and **OK**.

Modifying an IP alias interface

After you create an IP alias, you can update the settings to respond to network changes or resolve issues.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. In the **Interfaces** column, select the checkbox of the interface and disable the IP alias interface by clicking **No** in the **Enabled** column. Click **OK** in the warning dialog box.
3. In the **Interfaces** column, select the checkbox of the interface and click **Configure**.
4. In the Configure IP Alias dialog box, change the settings as described in the procedure for creating an IP Alias.
5. Click **Next** and **Finish**.

Registering interfaces with DDNS

Dynamic DNS (DDNS) is a protocol that registers local IP addresses on a Domain Name System (DNS) server.

About this task

In this release, DD System Manager supports Windows mode DDNS. To use UNIX mode DDNS, use the `net ddns` CLI command. You can do the following.

- Manually register (add) configured interfaces to the DDNS registration list.
- Remove interfaces from the DDNS registration list.
- Enable or disable DNS updates.
- Display whether DDNS registration is enabled or not.
- Display interfaces in the DDNS registration list.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces** > **DDNS Registration**.
2. In the DDNS Windows Mode Registration dialog box, click **Add** to add an interface to the DDNS.

   The Add Interface dialog box appears.

   a. Enter a name in the **Interface** field.
   b. Click **OK**.
3. Optionally, to remove an interface from the DDNS:
   a. Select the interface to remove, and click **Remove**.
   b. In the Confirm Remove dialog box, click **OK**.
4. Specify the DDNS Status.
   - Select **Enable** to enable updates for all interfaces already registered.
   - Click **Default** to select the default settings for DDNS updates.
   - Clear **Enable** to disable DDNS updates for the registered interfaces.
5. To complete the DDNS registration, click **OK**.

Destroying an interface

You can use DD System Manager to destroy or delete bonded, VLAN, and IP alias interfaces.

About this task

When a bonded interface is destroyed, the system deletes the bonded interface, releases its bonded physical interface, and deletes any VLANs or aliases attached to the bonded interface. When you delete a VLAN interface, the OS deletes the VLAN and any IP alias interfaces that are created under it. When you destroy an IP alias, the OS deletes only that alias interface.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces**.
2. Click the box next to each interface you want to destroy (Virtual or VLAN or IP Alias).
3. Click **Destroy**.
4. Click **OK** to confirm.

Viewing an interface hierarchy in the tree view

The Tree View dialog box displays the association between physical and bonded interfaces.

Steps

1. Select **Hardware** > **Ethernet** > **Interfaces** > **Tree View**.
2. In the Tree View dialog box, click the plus or minus boxes to expand or contract the tree view that shows the hierarchy.
3. Click **Close** to exit this view.

General network settings management

The configuration settings for hostname, domain name, search domains, host mapping, and DNS list are managed together on the Settings tab.

Viewing network settings information

The Settings tab displays the current configuration for the hostname, domain name, search domains, host mapping, and DNS.

Steps

- Select **Hardware** > **Ethernet** > **Settings**.

Setting the DD System Manager hostname

You can configure the DD System Manager hostname and domain name manually, or you can configure DDOS to automatically receive the host and domain names from a Dynamic Host Configuration Protocol (DHCP) server.

About this task

One advantage to manually configuring the host and domain names is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, if possible, manually configure the host and domain names.

When configuring the hostname and domain name, consider the following guidelines.

- Do not include an underscore in the hostname; it is incompatible with some browsers.
- Replication and CIFS authentication must be reconfigured after you change the names.
- If a system was previously added without a fully qualified name (no domain name), a domain name change requires that you remove and add the affected system or update the Search Domain List to include the new domain name.

NOTE: For steps on changing an established hostname, see the KB article *Data Domain - Changing Data Domain Hostname*, available from the Online Support website.

Steps

1. Select **Hardware** > **Ethernet** > **Settings**.
2. Click **Edit** in the **Host Settings** area. The **Configure Host** dialog opens.
3. To manually configure the host and domain names:
   a. Select **Manually configure host**.
   b. Enter a hostname in the **Host Name** box.

      For example, **`id##.yourcompany.com`**
   c. Enter a domain name in the **Domain Name** box.

      This is the domain name that is associated with your protection system and, usually, your company's domain name. For example, *yourcompany.com*
   d. Click **OK**.

      The system displays progress messages as the changes are applied.
4. To obtain the host and domain names from a DHCP server, select **Obtain Settings using DHCP** and click **OK**.

   At least one interface must be configured to use DHCP.

Managing the domain search list

Use the domain search list to define which domains the system can search.

Steps

1. Select **Hardware** > **Ethernet** > **Settings**.
2. Click **Edit** in the Search Domain List area.
3. To add a search domain using the Configure Search Domains dialog:
   a. Click Add (**+**).
   b. In the Add Search Domain dialog, enter a name in the **Search Domain** box.

      For example, **`id##.yourcompany.com`**
   c. Click **OK**.

      The system adds the new domain to the list of searchable domains.
   d. Click **OK** to apply changes and return to the Settings view.
4. To remove a search domain using the Configure Search Domains dialog:
   a. Select the search domain to remove.
   b. Click Delete (**X**).

      The system removes the selected domain from the list of searchable domains.
   c. Click **OK** to apply changes and return to the Settings view.

Adding and deleting host maps

A host map links an IP address to a hostname, so that either the IP address or the hostname can be used to specify the host.

Steps

1. Select **Hardware** > **Ethernet** > **Settings**.
2. To add a host map, do the following.
   a. In the Hosts Mapping area, click **Add**.
   b. In the Add Hosts dialog, enter the IP address of the host in the **IP Address** box.
   c. Click Add (**+**).
   d. In the Add Host dialog, enter a hostname, such as **`id##.yourcompany.com`**, in the **Host Name** box.
   e. Click **OK** to add the new hostname to the Host Name list.

f. Click **OK** to return to the Settings tab.

3. To delete a host map, do the following.
   a. In the Hosts Mapping area, select the host mapping to delete.
   b. Click Delete (**X**).

Configuring DNS IP addresses

DNS IP addresses specify the DNS servers the system can use to get IP addresses for host names that are not in the host mapping table.

About this task

You can configure the DNS IP addresses manually, or you can configure DDOS to automatically receive IP addresses from a DHCP server. One advantage to manually configuring DNS IP addresses is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, Dell recommends that you manually configure the DNS IP addresses.

Steps

1. Select **Hardware** > **Ethernet** > **Settings**.
2. Click **Edit** in the DNS List area.
3. To manually add a DNS IP address:
   a. Select **Manually configure DNS list**.

      The DNS IP address checkboxes become active.
   b. Click Add (**+**).
   c. In the Add DNS dialog box, enter the DNS IP address to add.
   d. Click **OK**.

      The system adds the new IP address to the list of DNS IP addresses.
   e. Click **OK** to apply the changes.
4. To delete a DNS IP address from the list:
   a. Select **Manually configure DNS list**.

      The DNS IP address checkboxes become active.
   b. Select the DNS IP address to delete and click Delete (**X**).

      The system removes the IP address from the list of DNS IP addresses.
   c. Click **OK** to apply the changes.
5. To obtain DNS addresses from a DHCP server, select **Obtain DNS using DHCP** and click **OK**.

   At least one interface must be configured to use DHCP.

Network route management

Routes determine the path taken to transfer data to and from the localhost (the protection system) to another network or host.

Data Domain and PowerProtect systems do not generate or respond to any of the network routing management protocols (RIP, EGRP/EIGRP, and BGP). The only routing implemented on a protection system is IPv4 policy-based routing, which allows only one route to a default gateway per routing table. There can be multiple routing tables and multiple default gateways. A routing table is created for each address that has the same subnet as a default gateway. The routing rules send the packets with the source IP address that matches the IP address used to create the table to that routing table. All other packets that do not have source IP addresses that match a routing table are sent to the main routing table.

Within each routing table, static routes can be added, but because source routing is used to get packets to the table, the only static routes that will work are static routes that use the interface that has the source address of each table. Otherwise it needs to be put into the main table.

Static routes are also required in the main routing table to direct which source addresses to use with connections initiated from DDOS if the destination program does not bind the IP address.

NOTE: DD Replicator sets a static route between the source and target systems when the replication context is created, therefore it does not require the creation of additional static routes.

Other than the IPv4 source routing done to these other routing tables, Data Domain and PowerProtect systems use source-based routing for the main routing IPv4 and IPv6 tables, which means that outbound network packets that match the subnet of multiple interfaces are routed only over the physical interface whose IP address matches the source IP address of the packets, which is where they originated.

For IPv6, set static routes when multiple interfaces contain the same IPv6 subnets, and the connections are being made to IPv6 addresses with this subnet. Normally, static routes are not needed with IPv4 addresses with the same subnet, such as for backups. There are cases in which static addresses may be required to allow connections to work, such as connections from the protection system to remote systems.

Static routes can be added and deleted from individual routing tables by adding or deleting the table from the route specification. This provides the rules to direct packets with specific source addresses through specific route tables. If a static route is required for packets with those source addresses, the routes must be added the specific table where the IP address is routed.

NOTE: Routing for connections initiated from the protection system, such as for replication, depends on the source address used for interfaces on the same subnet. To force traffic for a specific interface to a specific destination (even if that interface is on the same subnet as other interfaces), configure a static routing entry between the two systems: this static routing overrides source routing. This is not needed if the source address is IPv4 and has a default gateway associated with it. In that case, the source routing is already handled via its own routing table.

Viewing route information

The Routes tab displays the default gateways, static routes, and dynamic routes.

Steps

Select **Hardware** > **Ethernet** > **Routes**.

NOTE: If this does not display all the routing tables configured on the system, run the `net route show tables` command to display all the tables. The *DDOS Command Reference Guide* provides additional information.

Results

The Static Routes area lists the route specification used to configure each static route. The Dynamic Routes table lists information for each of the dynamically assigned routes.

Setting the default gateway

You can configure the default gateway manually, or you can configure DDOS to automatically receive the default gateway IP addresses from a DHCP server.

About this task

One advantage to manually configuring the default gateway is that you remove the dependency on the DHCP server and the interface leading to the DHCP server. To minimize the risk of service interruption, if possible, manually configure the default gateway IP address.

NOTE: The system supports the use of additional default gateways that are configured on specific NICs. Use the `net route add gateway` command to configure additional default gateways. The *DDOS Command Reference Guide* provides additional information.

Steps

1. Select **Hardware** > **Ethernet** > **Routes**.
2. Click **Edit** next to the default gateway type (IPv4 or IPv6) you want to configure.
3. To manually configure the default gateway address:
   a. Select **Manually Configure**.
   b. Enter the gateway address in the **Gateway** box.
   c. Click **OK**.
4. To obtain the default gateway address from a DHCP server, select **Use DHCP value** and click **OK**.

At least one interface must be configured to use DHCP.

Creating static routes

Static routes define destination hosts or networks that they system can communicate with.

About this task

NOTE: The steps for adding a static route using the CLI can be found in the KB article, *Data Domain: How to add static route to go through a specific interface on DDOS 6.0?*, available from the Online Support website.

Steps

1. Select **Hardware** > **Ethernet** > **Routes**.
2. Click **Create** in the Static Routes area.
3. In the **Create Routes** dialog, select the interface you want to host the static route, and click **Next**.
4. Specify the Destination.
   - To specify a destination network, select **Network** and enter the network address and netmask for the destination network.
   - To specify a destination host, select **Host** and enter the hostname or IP address of the destination host.
5. Optionally, specify the gateway to use to connect to the destination network or host.
   a. Select **Specify a gateway for this route**.
   b. Enter the gateway address in the **Gateway** box.
6. Review the configuration and click **Next**.

   The create routes Summary page appears.
7. Click **Finish**.
8. After the process is completed, click **OK**.

   The new route specification is listed in the Route Spec list.

Deleting static routes

Delete a static route when you no longer want the system to communicate with a destination host or network.

Steps

1. Select **Hardware** > **Ethernet** > **Routes**.
2. Select the Route Spec of the route specification to delete.
3. Click **Delete**.
4. Click **Delete** to confirm and then click **Close**.

   The selected route specification is removed from the Route Spec list.

11

Managing Storage

This chapter includes:

Topics:

Managing DD system storage

Storage management features enable you to view the status and configuration of your storage space, flash a disk LED to facilitate disk identification, and change the storage configuration.

NOTE: All storage that is connected or used by the two-node Active-Standby HA system can be viewed as a single system.

Using the CLI to calculate usable storage space

The following values are required to calculate the usable storage on a protection system after accounting for RAID overhead:

- *N*= Number of disks in use in the disk group (dg).
- *C* = Capacity of each disk after formatting.
- *R* = 2 (Number of disks used for RAID 6 parity)

This calculation does not work for Cache Tier storage, because the Cache Tier disks are not RAID protected.

Run the `storage show all` command to get the values for *N* and *C*.

```
sysadmin@ddbeta90# storage show all
Active tier details:
Disk       Disks       Count   Disk       Additional
Group                          Size       Information
-------    --------    -----   --------   -----------
dg2        2.1-2.14    14      2.7 TiB
(spare)    2.15        1       2.7 TiB
-------    --------    -----   --------   -----------

Current active tier size: 32.7 TiB
Active tier maximum capacity: 131.0 TiB
```

Figure 3. Example of `storage show all` command

In this example there are 14 disks in use in dg2 and each disk has a capacity of 2.7 TiB, therefore *N*=14 and *C*= 2.7 TiB

Use the formula (*N*-*R*) x *C* to get the usable capacity. In this example, the equation is (14-2) x 2.7 TiB.

12 x 2.7 TiB = 32.4 TiB, or 35.6 TB.

NOTE: The calculated value may not match exactly with the output of the `storage show all` command due to the way the capacity values are rounded for display. The `disk show hardware` command displays the disk capacity with additional decimal places.

Viewing system storage information

The storage status area shows the current status of the storage, such as Operational or Non-Operational, and the storage migration status. Below the Status area are tabs that organize how the storage inventory is presented.

Steps

1. To display the storage status, select **Hardware** > **Storage**.
2. If an alerts link appears after the storage status, click the link to view the storage alerts.
3. If the Storage Migration Status is Not licensed, you can click **Add License** to add the license for this feature.

Physically locating an enclosure

If you have trouble determining which physical enclosure corresponds to an enclosure displayed in DD System Manager, you can use the CLI beacon feature to flash the enclosure IDENT LEDs and all the disk LEDs that indicate normal operation.

Steps

1. Establish a CLI session with the system.
2. Type `enclosure beacon` *enclosure*.
3. Press **`Ctrl-C`** to stop the LED flashing.

Physically locating a disk

If you have trouble determining which physical disk corresponds to a disk displayed in DD System Manager, you can use the beacon feature to flash an LED on the physical disk.

Steps

1. Select **Hardware** > **Storage** > **Disks**.
2. Select a disk from the **Disks** table and click **Beacon**.

   NOTE: You can select one disk at a time.

   NOTE: For DD6400 systems, the button is **Blink LED**.

   The Beaconing Disk dialog box appears, and the LED light on the disk begins flashing.
3. Click **Stop** to stop the LED beaconing.

Configuring storage

Storage configuration features allow you to add and remove storage expansion enclosures from the active, retention, and cloud tiers. Storage in an expansion enclosure (also called an expansion shelf) is not available for use until it is added to a tier.

About this task

NOTE: Additional storage requires the appropriate license or licenses and sufficient memory to support the new storage capacity. Error messages appear if more licenses or memory is needed.

DD6300 systems support the option to use enclosures with 4 TB drives (43.6 TiB) at 50% utilization (21.8 TiB) in the active tier.

The following guidelines apply to using partial capacity shelves with DD6300 systems:

- The available licensed capacity must be exactly 21.8 TiB.
- No other enclosure types or drive sizes are supported for use at partial capacity.
- A partial shelf can only exist in the Active tier.
- Only one partial ES30 can exist in the Active tier.
- Once a partial shelf exists, no additional ES30s can be configured until the partial shelf is added at full capacity. You must license enough additional capacity to use the remaining 21.8 TiB of the partial shelf.
- If the available capacity exceeds 21.8 TB, you cannot add a partial shelf.
- Deleting a 21 TiB license does not automatically convert a fully used shelf to a partial shelf. You must remove the shelf and add it back as a partial shelf.

DD6900 systems support the option to use enclosures with 4 TB drives (43.6 TiB) at 25% utilization (10.9 TiB), 50% utilization (21.8 TiB), or 75% (32.7 TiB) in the active tier.

The following guidelines apply to using partial capacity shelves with DD6900 systems:

- DS60 disk packs and ES40 shelves must be fully licensed before adding another partial shelf license.
- A partial shelf can only exist in the Active tier.
- Only one partial shelf can exist in the Active tier.
- Once a partial shelf exists, no additional partial shelves or packs can be configured until the partial shelf is added at full capacity. You must license enough additional capacity to use the remaining capacity of the partial shelf.
- Controller upgrades from a DD6900 system cannot be completed while a partial shelf exists on the system.
- Storage migration cannot be completed while a partial shelf exists on the system.

For DD6900, DD9400, and DD9900 systems, storage capacity licenses are available in increments of 60 TB raw (48 TB usable) capacity. For systems with 8 TB drives, the licensed capacity might not equal the full capacity of the disks that are installed in the disk shelves. For example, if the system has one pack of 8 TB disks (96 TB usable capacity) but the licensed capacity is 48 TB, only half the system capacity is available for use.

For DD6400 systems, go to the *Configuring Storage for DD6400 systems section*.

Steps

1. Select **Hardware** > **Storage** > **Overview**.
2. Expand the dialog box for the **Active Tier**, **Cache Tier**, or **Cloud Tier**.
3. Click **Configure**.
4. In the **Configure Storage** dialog box, select the storage from the **Addable Storage** list.
5. In the **Configure** list, select **Active Tier**.

   The maximum amount of storage that can be added to the active tier depends on the active DD controller.

   NOTE: The licensed capacity bar shows the portion of licensed capacity (used and remaining) for the installed enclosures.

6. Select the shelf and click **Add to Tier**.
7. Click **OK** to add the storage.

   NOTE: To remove an added shelf, select it in the Tier Configuration list, click **Remove from Configuration**, and click **OK**.

Configuring storage for DD6400 systems

Prerequisites

Administrator privileges are required to configure storage on a DD6400 system.

About this task

DD6400 systems use a different method to configure storage than other DD systems. For a DD6400, the storage is configured automatically based on the information in the Protection Pool license.

Steps

If storage was not configured on Initial Configuration or additional capacity has been added, click **Add Storage**.

Expanding DD3300 capacity

The DD3300 system is available in four different capacity configurations. You can expand capacity from one configuration to another.

Prerequisites

All capacity expansions require the installation of additional disks and memory in the system. Do not attempt to expand the capacity until the hardware upgrades are complete.

About this task

The DD3300 system is available in 4 TB, 8 TB, 16 TB, and 32 TB capacity configurations. Capacity expansion is a one-time process.

The following table shows the available expansion paths and required hardware:

Table 7. DD3300 upgrade requirements for capacity expansion

Capacity expansion	Additional memory	Additional HDDs	Additional SSD
4 TB to 16 TB	32 GB	6 x 4 TB HDDs	1 x 480 GB SSD
8 TB to 16 TB	8 TB to 16 TB expansion requires licensing and configuration changes only. No hardware upgrades are required.		
16 TB to 32 TB	16 GB	6 x 4 TB HDDs	N/A

Steps

1. select **Maintenance** > **System**.

   The **Capactiy Expansion History** table displays details about the capacity of the system, including the initial capacity of the system, the date of the initial software installation, and the expanded capacity and the date of the expansion operation if applicable.

2. If the system has not been expanded, select the target capacity from the **Select Capacity** list box and click **Capacity Expand** to initiate the capacity expansion.

   The *DD3300 Field Replacement and Upgrade Guide* provides detailed instructions for expanding system capacity.

   NOTE: Insufficient memory, insufficient physical capacity (HDDs), the system has already been expanded, or the target for capacity expansion is not supported can prevent a capacity expansion. If the capacity expansion cannot be completed, the system displays the reason.

Failing and unfailing a disk

Disk fail functionality allows you to manually set a disk to a failed state to force reconstruction of the data stored on the disk. Disk unfail functionality allows you to take a disk in a failed state and return it to operation.

About this task

DD6400 systems use spare extents spread across all the disks in the tier rather than dedicated spare disks. Therefore, the estimated time for disk reconstruction is calculated by analyzing the available spare extents and providing an estimate based on the amount of spare space available. If additional storage is added while reconstruction is ongoing, the reconstruction time will increase because there is more spare space available. If storage is added after reconstruction stops due to insufficient space, reconstruction resumes with the new space

Steps

- To fail a disk and force reconstruction:

1. Select **Hardware** > **Storage** > **Disks** > **Fail**.
2. Select a disk from the table and click **Fail**.

   NOTE: For DD6400 systems, the button is **Remove**.

- To make a disk marked Failed or Foreign usable to the system:
  1. Select **Hardware** > **Storage** > **Disks** > **Unfail**.
  2. Select a disk from the table and click **Unfail**.

     NOTE: For DD6400 systems, the button is **Add**.

12

Migrating Storage

This chapter presents the following topics:

Topics:

Storage migration overview

Storage migration supports the replacement of existing storage enclosures with new enclosures that may offer higher performance, higher capacity, and a smaller footprint.

After new enclosures are installed, you can migrate the data from the older enclosures to the new enclosures while the system continues to support other processes such as data access, expansion, cleaning, and replication. The storage migration does require system resources, but you can control this with throttle settings that give the migration a relatively higher or lower priority. You can also suspend a migration to make more resources available to other processes, then resume the migration when resource demand is lower.

During the migration, the system uses data on the source and destination enclosures. New data is written to the new enclosures. Non-migrated data is updated on the source enclosures, and migrated data is updated on the destination enclosures. If the migration is interrupted, the migration can resume migrating blocks that have not been marked as migrated.

During the migration, each block of data is copied and verified, the source block is freed and marked as migrated, and the system index is updated to use the new location. New data that was destined to land in the source block will now be redirected to destination block. All new data block allocations that would have been allocated from source are allocated from the destination.

The Migration copy process is done at the shelf level, not the logical data level, so all disk sectors on the source shelf are accessed and copied over regardless of whether there is data on them. Therefore, the Storage Migration Utility cannot be used to shrink a logical data footprint.

NOTE: Because the data set is divided between the source and destination enclosures during migration, you cannot halt a migration and resume use of only the source enclosures. Once started, the migration must complete. If a failure, such as a faulty disk drive, interrupts the migration, address the issue and resume the migration.

Depending on the amount of data to migrate and the throttle settings selected, a storage migration can take days or weeks. When all data is migrated, the finalize process, which must be manually initiated using the `storage migration finalize` command, restarts the file system. During the restart, the source enclosures are removed from the system configuration and the destination enclosures become part of the file system. When the finalize process is complete, the source enclosures can be removed from the system.

After a storage migration, the disk shelf numbers reported by DDOS might not be sequential. This is because shelf numbering is tied to the serial number of each individual disk shelf. The KB article *Data Domain: Storage enclosure numbering is not sequential*, available from the Online Support website, provides additional details. In DDOS version 5.7.3.0 and later, the `enclosure show persistent-id` command described in the KB article requires administrator access, not SE access.

Migration planning considerations

Consider the following guidelines before starting a storage migration.

- Storage migration requires a single-use license and operates on system models supported by DDOS version 5.7 or later.
  NOTE: Multiple storage migration operations require multiple licenses. However, multiple source enclosures can be migrated to multiple destination enclosures during a single operation.
- All source shelves must be fully licensed. The migration cannot be completed while a partial shelf exists on the system.
- Two licenses are required for storage migration:
  - The storage migration feature license
  - The capacity and shelf type license for the destination enclosures
- Storage migration is based on capacity, not enclosure count. Therefore:
  - One source enclosure can be migrated to one destination enclosure.
  - One source enclosure can be migrated to multiple destination enclosures.
  - Multiple source enclosures can be migrated to one destination enclosure.
  - Multiple source enclosures can be migrated to multiple destination enclosures.
- The storage migration licensing process consists of:
  1. Adding the storage migration feature license and appending the new storage capacity and shelf type to the existing licensed shelf capacity before running the migration operation.
     NOTE: This may result in the licensed storage capacity exceeding the maximum capacity supported by the DD model for the duration of the storage migration operation. Without this, the storage migration cannot start.
  2. Removing the original capacity and shelf type license and the storage migration feature license after the migration operation is complete, leaving only the new capacity and shelf type license.
- The destination enclosures must:
  - Be unassigned shelves with the drives in an unused state.
  - Be licensed for sufficient capacity to receive the data from the source enclosures, with the license installed on the system
  - Be supported on the DD system model.
  - Contain at least as much usable capacity as the enclosures they are replacing.
    NOTE: It is not possible to determine the utilization of the source shelf. The system performs all calculations based on the capacity of the shelf.
- The DD system model must have sufficient memory to support the active tier storage capacity of the new enclosures.
- Data migration is not supported for disks in the system controller.
- **CAUTION: Do not upgrade DDOS until the in-progress storage migration is complete.**
- Storage migration cannot start when the file system is disabled or while a DDOS upgrade is in progress, another migration is in progress, or a RAID reconstruction is in progress.
  NOTE: If a storage migration is in progress, a new storage migration license is required to start a new storage migration operation after the in-progress migration completes. The presence or absence of a storage migration license is reported as part of the upgrade precheck.
- All specified source enclosures must be in the same tier (active).
- There can be only one disk group in each source enclosure, and all disks in the disk group must be installed in within the same enclosure.
- All disks in each destination enclosure must be of the same type (for example, all SATA or all SAS).
- After migration begins, the destination enclosures cannot be removed.
- Source enclosures cannot be removed until migration is complete and finalized.
- The storage migration duration depends on the system resources (which differ for different system models), the availability of system resources, and the data quantity to migrate. Storage migration can take days or weeks to complete.

8 TB shelf migration considerations

For systems that have been upgraded to a DD9400 or a DD9900, 3 TB and 4 TB shelves may be migrated to 8 TB shelves.

CAUTION: DDOS 7.4 or higher is required for 8 TB shelf migration. Do not attempt the 8 TB shelf migration on a system running DDOS 7.3 or lower.

For migration from 3 TB or 4 TB SAS drives to 8 TB SAS drives as a part of DD9400 or DD9900 controller upgrade with additional capacity, two additional 8 TB disk packs are required for overprovisioning if the additional capacity is greater than or equal to 192 TB. These drives are required for adequate performance of the 8 TB drives, and are provided as part of the controller upgrade with additional capacity only. The overprovisioning drives cannot be added later.

The following guidelines apply for migrations to 8 TB shelves:

- For systems running DDOS 7.4 or higher that were upgraded to a DD9400 or DD9900, a minimum capacity increase of 192 TB is required to add 8 TB drives to the system. The new capacity will also come with two disk packs for overprovisioning to meet performance targets.

  NOTE: Capacity increases lower than 192 TB are fulfilled with 4 TB shelves.
- When migrating from 4 TB shelves to 8 TB shelves on a system with overprovisioning shelves, the migration requires additional capacity beyond what is needed for migration because of the overprovisioning percentage.
- For systems running DDOS 7.3 or lower that were upgraded to a DD9400 or DD9900, overprovisioning shelves are not included with the new storage.

DS60 shelf considerations

The DS60 dense shelves can hold 60 disks, allowing the customer to use the full amount of space in the rack. The drives are accessed from the top of the shelf, by extending the shelf from the cabinet. Due to the weight of the shelves, approximately 225 lbs when fully loaded, read this section before proceeding with a storage migration to DS60 shelves.

Be aware of the following considerations when working with the DS60 shelf:

CAUTION:

- **Loading shelves at the top of the rack may cause the shelf to tip over.**
- **Validate that the floor can support the total weight of the DS60 shelves.**
- **Validate that the racks can provide enough power to the DS60 shelves.**
- **When adding more than five DS60s in the first rack, or more than six DS60s in the second rack, stabilizer bars and a ladder are required to maintain the DS60 shelves.**

Viewing migration status

DD System Manager provides two ways to view storage migration status.

Steps

1. Select **Hardware** > **Storage**.

   In the Storage area, review the Storage Migration Status line. If the status is Not Licensed, you must add a license before using any storage migration features. If the storage migration license is installed, the status can be one of the following: None, Starting, Migrating, Paused by User, Paused by System, Copy Completed - Pending Finalization, Finalizing, Failed during Copy, or Failed during Finalize.
2. If a storage migration is in progress, click **View Storage Migration** to view the progress dialogs.

   NOTE: The migration status shows the percentage of blocks transferred. In a system with many free blocks, the free blocks are not migrated, but they are included in the progress indication. In this situation, the progress indication will climb quickly and then slow when the data migration starts.
3. When a storage migration is in progress, you can also view the status by selecting **Health** > **Jobs**.

Evaluating migration readiness

You can use the system to evaluate storage migration readiness without committing to start the migration.

Steps

1. Install the destination enclosures using the instructions in the product installation guides.
2. Select **Administration** > **Licenses** and verify that the storage migration license is installed.
3. If the storage migration license is not installed, click **Add Licenses** and add the license.
4. Select **Hardware** > **Storage**, then click **Migrate Data**.
5. In the Select a Task dialog, select **Estimate**, then click **Next**.
6. In the Select Existing Enclosures dialog, use the checkboxes to select each of the source enclosures for the storage migration, then click **Next**.

7. In the Select New Enclosures dialog, use the checkboxes to select each of the destination enclosures for the storage migration, then click **Next**.

   The Add Licenses button allows you to add storage licenses for the new enclosures as needed, without interrupting the current task.

8. In the Review Migration Plan dialog, review the estimated migration schedule, then click **Next**.
9. Review the precheck results in the Verify Migration Preconditions dialog, then click Close.

Results

If any of the precheck tests fail, resolve the issue before you start the migration.

Migrating storage using DD System Manager

The storage migration process evaluates system readiness, prompts you to confirm that you want to start the migration, migrates the data, and then prompts you to finalize the process.

Steps

1. Install the destination enclosures using the instructions in the product installation guides.
2. Select **Administration** > **Licenses** and verify that the storage migration license is installed.
3. If the storage migration license is not installed, click **Add Licenses** and add the license.
4. Select **Hardware** > **Storage**, then click **Migrate Data**.
5. In the Select a Task dialog, select **Migrate**, then click **Next**.
6. In the Select Existing Enclosures dialog, use the checkboxes to select each of the source enclosures for the storage migration, then click **Next**.
7. In the Select New Enclosures dialog, use the checkboxes to select each of the destination enclosures for the storage migration, then click **Next**.

   The Add Licenses button allows you to add storage licenses for the new enclosures as needed, without interrupting the current task.

8. In the Review Migration Plan dialog, review the estimated migration schedule, then click **Start**.
9. In the Start Migration dialog, click **Start**.

   The Migrate dialog appears and updates during the three phases of the migration: Starting Migration, Migration in Progress, and Copy Complete.

10. When the Migrate dialog title displays Copy Complete and a filesystem restart is acceptable, click **Finalize**.

    NOTE: This task restarts the filesystem and typically takes 10 to 15 minutes. The system is unavailable during this time.

Results

When the migration finalize task is complete, the system is using the destination enclosures and the source enclosures can be removed.

Storage migration dialog descriptions

The DD System Manager dialog descriptions provide additional information on storage migration. This information is also available by clicking the help icon in the dialogs.

Select a Task dialog

The configuration in this dialog determines whether the system will evaluate storage migration readiness and stop, or evaluate readiness and begin storage migration.

Select **Estimate** to evaluate system readiness and stop.

Select **Migrate** to start migration after the system evaluation. Between the system evaluation and the start of the migration, a dialog prompts you to confirm or cancel the storage migration.

Select Existing Enclosures dialog

The configuration in this dialog selects either the active or the retention tier and the source enclosures for the migration.

The Existing Enclosures list displays the enclosures that are eligible for storage migration. Select the checkbox for each of the enclosures to migrate. Click **Next** when you are ready to continue.

Select New Enclosures dialog

The configuration in this dialog selects the destination enclosures for the migration. This dialog also displays the storage license status and an **Add Licenses** button.

The Available Enclosures list displays the enclosures that are eligible destinations for storage migration. Select the checkbox for each of the desired destination enclosures.

The license status bar represents all of the storage licenses installed on the system. The green portion represents licenses that are in use, and the and clear portion represents the licensed storage capacity available for destination enclosures. If you need to install additional licenses to support the selected destination controllers, click **Add Licenses**.

Click **Next** when you are ready to continue.

Review Migration Plan dialog

This dialog presents an estimate of the storage migration duration, organized according to the three stages of storage migration.

Stage 1 of the storage migration runs a series of tests to verify that the system is ready for the migration. The test results appear in the Verify Migration Preconditions dialog.

During Stage 2, the data is copied from the source enclosures to the destination enclosures. When a large amount of data is present, the copy can take days or weeks to complete because the copy takes place in the background, while the system continues to serve backup clients. A setting in the Migration in Progress dialog allows you to change the migration priority, which can speed up or slow down the migration.

Stage 3, which is manually initiated from the Copy Complete dialog, updates the system configuration to use the destination enclosures and removes the configuration for the source controllers. During this stage, the file system is restarted and the system is unavailable to backup clients.

Verify Migration Preconditions dialog

This dialog displays the results of the tests that execute before the migration starts.

The following list shows the test sequence and provides additional information on each of the tests.

P1. This system's platform is supported.	Older DD system models do not support storage migration.
P2. A storage migration license is available.	A storage migration license is required.
P3. No other migration is currently running.	A previous storage migration must complete before you can start another.
P4. The current migration request is the same as	Resume and complete the interrupted migration.

the interrupted migration request.	
P5. Check the disk group layout on the existing enclosures.	Storage migration requires that each source enclosure contain only one disk group, and all the disks in the group must be in that enclosure.
P6. Verify the final system capacity.	The total system capacity after migration and the removal of the source enclosures must not exceed the capacity supported by the DD system model.
P7. Verify the replacement enclosures' capacity.	The usable capacity of the destination enclosures must be greater than that of the source enclosures.
P8. Source enclosures are in the same active tier or retention unit.	The system supports storage migration from either the active tier or the retention tier. It does not support migration of data from both tiers at the same time.
P9. Source enclosures are not part of the head unit.	Although the system controller is listed as an enclosure in the CLI, storage migration does not support migration from disks installed in the system controller.
P10. Replacement enclosures are addable to storage.	All disks in each destination enclosure must be of the same type (for example, all SATA or all SAS), with the capacity and shelf type license installed on the system.
P11. No RAID reconstruction is occurring in the source controllers.	Storage migration cannot start while a RAID reconstruction is in progress.
P12. Source shelf belongs to a supported tier.	The source disk enclosure must be part of a tier supported on the migration destination.

Migration progress dialogs

This series of dialogs presents the storage migration status and the controls that apply at each stage.

Migrate - Starting Migration

During the first stage, the progress is shown on the progress bar and no controls are available.

Migrate - Migration in Progress

During the second stage, data is copied from the source enclosures to the destination enclosures and the progress is shown on the progress bar. Because the data copy can take days or weeks to complete, controls are provided so that you can manage the resources used during migration and suspend migration when resources are needed for other processes.

You can click **Pause** to suspend the migration and later click **Resume** to continue the migration.

The **Low**, **Medium**, and **High** buttons define throttle settings for storage migration resource demands. A low throttle setting gives storage migration a lower resource priority, which results in a slower migration and requires fewer system resources. Conversely, A high throttle setting gives storage migration a higher resource priority, which results in a faster migration and requires more system resources. The medium setting selects an intermediate priority.

You do not have to leave this dialog open for the duration of the migration. To check the status of the migration after closing this dialog, select **Hardware** > **Storage** and view the migration status. To return to this dialog from the Hardware/Storage page, click **Manage Migration**. The migration progress can also be viewed by selecting **Health** > **Jobs**.

Migrate - Copy Complete

When the copy is complete, the migration process waits for you to click **Finalize**. During this final stage, , which takes 10 to 15 minutes, the filesystem is restarted and the system is not available. It is a good practice to start this stage during a maintenance window or a period of low system activity.

Migrating storage using the CLI

About this task

A migration simply requires moving all of the allocated blocks from the blocksets formatted over source DGs (e.g., source blocksets) to the blocksets formatted over destination DGs (e.g., destination blocksets). Once all of the allocated blocks have been moved from the source blocksets, those blocksets can be removed from the file system, their disks can be removed from their storage tier, and the physical disks and enclosures can be removed from the DDR.

NOTE: The preparation of new enclosures for storage migration is managed by the storage migration process. Do not prepare destination enclosures as you would for an enclosure addition. For example, use of the `filesys expand` command is appropriate for an enclosure addition, but this command prevents enclosures from being used as storage migration destinations.

A DS60 disk shelf contains four disk packs, of 15 disks each. When a DS60 shelf is the migration source or destination, the disk packs are referenced as ***enclosure:pack***. In this example, the source is enclosure 7, pack 2 (7:2), and the destination is enclosure 7, pack 4 (7:4).

Steps

1. Install the destination enclosures using the instructions in the product installation guides.
2. Check to see if the storage migration feature license is installed.

   ```
   # elicense show
   ```
3. If the license is not installed, update the elicense to add the storage migration feature license

   ```
   # elicense update
   ```
4. View the disk states for the source and destination disks.

   ```
   # disk show state
   ```

   The source disks should be in the active state, and the destination disks should be in the unknown state.
5. Run the storage migration precheck command to determine if the system is ready for the migration.

   ```
   # storage migration precheck source-enclosures 7:2 destination-enclosures 7:4
   ```
6. View the migration throttle setting.

   ```
   storage migration option show throttle
   ```
7. When the system is ready, begin the storage migration.

   ```
   # storage migration start source-enclosures 7:2 destination-enclosures 7:4
   ```
8. Optionally, view the disk states for the source and destination disks during the migration.

   ```
   # disk show state
   ```

   During the migration, the source disks should be in the migrating state, and the destination disks should be in the destination state.
9. Review the migration status as needed.

   ```
   # storage migration status
   ```

10. View the disk states for the source and destination disks.

```
# disk show state
```

During the migration, the source disks should be in the migrating state, and the destination disks should be in the destination state.

11. When the migration is complete, update the configuration to use the destination enclosures.

NOTE: This task restarts the file system and typically takes 10 to 15 minutes. The system is unavailable during this time.

```
storage migration finalize
```

12. If you want to remove all data from each of the source enclosures, remove the data now.

```
storage sanitize start enclosure <enclosure-id>[:<pack-id>]
```

NOTE: The storage sanitize command does not produce a certified data erasure. Dell offers certified data erasure as a service. For more information, contact your Dell representative.

13. View the disk states for the source and destination disks.

```
# disk show state
```

After the migration, the source disks should be in the unknown state, and the destination disks should be in the active state.

Results

When the migration finalize task is complete, the system is using the destination storage and the source storage can be removed.

CLI storage migration example

elicense show

```
# elicense show
Feature licenses:
## Feature     Count Mode            Expiration Date
-- ----------- ----- --------------- ---------------
1  REPLICATION 1     permanent (int) n/a
2  VTL         1     permanent (int) n/a
-- ----------- ----- --------------- ---------------
```

elicense update

```
# elicense update mylicense.lic
New licenses: Storage Migration
Feature licenses:
## Feature           Count   Mode              Expiration Date
-- -----------        -----   ---------------   ---------------
1  REPLICATION        1       permanent (int)   n/a
2  VTL                1       permanent (int)   n/a
3  Storage Migration  1       permanent (int)
-- -----------        -----   ---------------   ---------------
** This will replace all existing Data Domain licenses on the system with the above Dell
ELMS licenses.
Do you want to proceed? (yes|no) [yes]: yes
eLicense(s) updated.
```

disk show state

```
# disk show state
Enclosure        Disk
  Row(disk-id)    1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
--------------   ---------------------------------------------
1                .  .  .  .
2                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
4                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
5                v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
7               |--------|--------|--------|--------|
                | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
     E(49-60)   |U  U  U |.  .  s |U  U  U |U  U  U |
     D(37-48)   |U  U  U |.  .  . |U  U  U |U  U  U |
     C(25-36)   |U  U  U |.  .  . |U  U  U |U  U  U |
     B(13-24)   |U  U  U |.  .  . |U  U  U |U  U  U |
     A( 1-12)   |U  U  U |.  .  . |U  U  U |U  U  U |
                |--------|--------|--------|--------|
--------------   ---------------------------------------------

Legend   State             Count
------   ---------------   -----
.        In Use Disks      18
s        Spare Disks       1
v        Available Disks   15
U        Unknown Disks     105
------   ---------------   -----
```

Figure 4. `disk show state`

storage migration precheck

```
#storage migration precheck  source-enclosures 2  destination-enclosures 11

  Source enclosures:
  Disks       Count   Disk     Disk          Enclosure   Enclosure
                      Group    Size          Model       Serial No.
  --------    -----   -----    ----------    ---------   --------------
  2.1-2.15    15      dg1      1.81 TiB      ES30        APM00111103820
  --------    -----   -----    ----------    ---------   --------------
  Total source disk size: 27.29 TiB

  Destination enclosures:
  Disks         Count   Disk       Disk        Enclosure   Enclosure
                        Group      Size        Model       Serial No.
  ----------    -----   -------    --------    ---------   --------------
  11.1-11.15    15      unknown    931.51 GiB  ES30        APM00111103840
  ----------    -----   -------    --------    ---------   --------------
  Total destination disk size: 13.64 TiB

       (P1)    Verifying platform support.....................................PASS
       (P2)    Verifying enclosure 1 is not used as source....................PASS
       (P3)    Verifying valid storage migration license exists...............PASS
       (P4)    Verifying no other migration is running........................PASS
       (P5)    Verifying data layout on the source shelves....................PASS
       (P6)    Verifying source shelves belong to same tier...................PASS
       (P7)    Verifying source shelves belong to a supported tier............PASS
       (P8)    Verifying 8TB shelf configuration..............................PASS
       (P9)    Verifying The number of shelves for performance reason.........PASS
       (P10)   Verifying final system capacity...............................PASS
       (P11)   Verifying destination capacity.................................PASS
       (P12)   Verifying destination shelves are addable to storage...........PASS
       (P13)   Verifying number of shelves after migration completes..........PASS
       (P14)   Verifying no RAID reconstruction is going on in source shelves....PASS
       (P15)   Verifying no RAID reconstruction is going on in 8TB shelves.....PASS
Migration pre-check PASSED

Expected time to migrate data: 8 hrs 33 min
```

storage migration show history

```
# storage migration show history
Id   Source        Source Enclosure   Dest          Dest Enclosure    Status      Start Time                 End Time
     Enclosure*    Serial No.         Enclosure*    Serial No.
--   ----------    ----------------   ----------    --------------    ---------   ------------------------   ------------------------
2    9:0           SHU952400106A23    7:0           SHU9524084G055B   Finalized   Sat Aug  8 11:59:37 2015   Mon Aug 10 11:10:11 2015
1    9:0           SHU952400106A23    7:0           SHU9524084G055B   Finalized   Thu Aug  6 16:39:55 2015   Fri Aug  7 10:28:07 2015
                                      8:0           SHU9524084G04LR
--   ----------    ----------------   ----------    --------------    ---------   ------------------------   ------------------------
(*) Enclosure ids at migration start time.
```

Figure 5. `storage migration show history`

storage migration start

```
#storage migration start  source-enclosures 2  destination-enclosures 11

  Source enclosures:
  Disks       Count   Disk     Disk          Enclosure   Enclosure
                      Group    Size          Model       Serial No.
  --------    -----   -----    ----------    ---------   --------------
  2.1-2.15    15      dg1      1.81 TiB      ES30        APM00111103820
  --------    -----   -----    ----------    ---------   --------------
  Total source disk size: 27.29 TiB

  Destination enclosures:
  Disks         Count   Disk       Disk        Enclosure   Enclosure
                        Group      Size        Model       Serial No.
  ----------    -----   -------    --------    ---------   --------------
  11.1-11.15    15      unknown    931.51 GiB  ES30        APM00111103840
  ----------    -----   -------    --------    ---------   --------------
  Total destination disk size: 13.64 TiB

  Expected time to migrate data: 84 hrs 40 min

**    Storage migration once started cannot be aborted.
      Existing data on the destination shelves will be overwritten.
         Do you want to continue with the migration? (yes|no) [no]: yes

Performing migration pre-check:
       (P1)    Verifying platform support.....................................PASS
       (P2)    Verifying enclosure 1 is not used as source....................PASS
       (P3)    Verifying valid storage migration license exists...............PASS
       (P4)    Verifying no other migration is running........................PASS
       (P5)    Verifying data layout on the source shelves....................PASS
       (P6)    Verifying source shelves belong to same tier...................PASS
       (P7)    Verifying source shelves belong to a supported tier............PASS
       (P8)    Verifying 8TB shelf configuration..............................PASS
       (P9)    Verifying The number of shelves for performance reason.........PASS
       (P10)   Verifying final system capacity...............................PASS
       (P11)   Verifying destination capacity.................................PASS
       (P12)   Verifying destination shelves are addable to storage...........PASS
       (P13)   Verifying number of shelves after migration completes..........PASS
       (P14)   Verifying no RAID reconstruction is going on in source shelves....PASS
       (P15)   Verifying no RAID reconstruction is going on in 8TB shelves.....PASS

Migration pre-check PASSED

Storage migration will reserve space in the filesystem to migrate data.
Space reservation may add up to an hour or more based on system resources.

Storage migration process initiated.
Check storage migration status to monitor progress.
```

storage migration status

```
# storage migration status
Id   Source          Destination     State       Percent    Estimated Time to Complete   Current Throttle
     Enclosure(s)    Enclosure(s)                Complete                                Setting
--   ------------    ------------    ---------   --------   --------------------------   ----------------
5    7:2             7:4             migrating   45%        30 hrs 18 mins               high
--   ------------    ------------    ---------   --------   --------------------------   ----------------
```

Figure 6. `storage migration status`

disk show state, migration in progress

```
# disk show state
Enclosure        Disk
  Row(disk-id)    1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
--------------   ------------------------------------------------
1                 .  .  .  .
2                 U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3                 U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
4                 U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
5                 v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6                 U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
7                |--------|--------|--------|--------|
                 | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
     E(49-60)    |U  U  U |m  m  s |U  U  U |s  d  d |
     D(37-48)    |U  U  U |m  m  m |U  U  U |d  d  d |
     C(25-36)    |U  U  U |m  m  m |U  U  U |d  d  d |
     B(13-24)    |U  U  U |m  m  m |U  U  U |d  d  d |
     A( 1-12)    |U  U  U |m  m  m |U  U  U |d  d  d |
                 |--------|--------|--------|--------|
--------------   ------------------------------------------------

Legend   State              Count
------   -----------------  -----
.        In Use Disks       4
s        Spare Disks        2
v        Available Disks    15
U        Unknown Disks      90
m        Migrating Disks    14
d        Destination Disks  14
```

Figure 7. `disk show state, migration in progress`

storage migration finalize

```
# storage migration finalize

Storage migration finalize restarts the filesystem.
This can take several minutes and the filesystem is unavailable until the operation completes.
        Do you want to continue? (yes|no) [no]: yes

Performing migration finalization pre-check:
(P1)    Verifying storage migration is ready for finalization....PASS
(P2)    Verifying there are no foreign disks.....................PASS
(P3)    Verifying data layout on the source shelves..............PASS

Migration finalization pre-check PASSED
Finalizing the storage migration with id 5:

Notifying filesystem to finalize migration...

Done.

Disabling the filesystem
Please wait.........
The filesystem is now disabled.
Removing source enclosures from filesystem...

Done.

Removing source enclosures from storage tier...

Done.

Enabling the filesystem
Please wait.....................
The filesystem is now enabled.
Storage migration with id 5 from enclosure(s) 7.2 to enclosure(s) 7.4 has been finalized.
```

Figure 8. `storage migration finalize`

disk show state, migration complete

```
# disk show state
Enclosure        Disk
  Row(disk-id)    1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
--------------   ----------------------------------------------
1                .  .  .  .
2                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
4                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
5                v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6                U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
7               |--------|--------|--------|--------|
                | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
     E(49-60)   |U  U  U |U  U  U |U  U  U |s  .  . |
     D(37-48)   |U  U  U |U  U  U |U  U  U |.  .  . |
     C(25-36)   |U  U  U |U  U  U |U  U  U |.  .  . |
     B(13-24)   |U  U  U |U  U  U |U  U  U |.  .  . |
     A( 1-12)   |U  U  U |U  U  U |U  U  U |.  .  . |
                |--------|--------|--------|--------|
--------------   ----------------------------------------------

Legend   State            Count
------   ---------------  -----
.        In Use Disks     18
s        Spare Disks      1
v        Available Disks  15
U        Unknown Disks    105
------   ---------------  -----
```

Figure 9. `disk show state, migration complete`

NOTE: Currently storage migration is only supported on the active node. Storage migration is not supported on the standby node of an HA cluster.

13

Managing File Systems

This chapter presents the following topics:

Topics:

Supported interfaces

The file system supports the following interfaces:

- NFS
- CIFS
- DD Boost
- DD VTL

File system limitations

Review the file system limitations.

- HA failover—Access to files may be interrupted for up to 10 minutes during failover on High Availability systems. (DD Boost and NFS require additional time.)
- Path name length—The maximum length of a full path name (including the characters in `/data/col1/backup`) is 1023 bytes. The maximum length of a symbolic link is also 1023 bytes.
- Battery—For systems that use NVRAM, the operating system creates a low battery alert if the battery charge falls below 80% capacity and the file system is disabled.
- Inodes—An NFS or CIFS client request causes a DD system to report a capacity of about two billion inodes (files and directories). Systems can exceed that number, but the reporting on the client might be incorrect.
- Number of files—EMC recommends storing a total of no more than 1 billion files on a system. This limitation applies to the combined number of files stored in both the Active and Cloud storage tiers. Storing a larger number of files can adversely affect the performance and the length of cleaning, and some processes, such as file system cleaning, may run much longer with a very large number of files. For example, the enumeration phase of cleaning may take from a few minutes to several hours depending upon the number of files in the system.

  NOTE: The overall system performance will fall to unacceptable levels if the system is required to support the maximum file amount and the workload from the client machines is not carefully controlled.

  When the file system passes the billion file limit, several processes or operations might be adversely affected, for example:

  - Cleaning may take a very long time to complete, perhaps several days.
  - AutoSupport operations may take more time.
  - Any process or command that needs to enumerate all the files.

  If there are many small files, other considerations arise:

  - The number of separate files that can be created per second, (even if the files are very small) may be more of a limitation than the number of MB/s that can be ingested. When files are large, the file creation rate is not significant, but when

files are small, the file creation rate dominates and may become a factor. The file creation rate is about 100 to 200 files per second depending upon the number of MTrees and CIFS connections. This rate should be taken into account during system sizing when a bulk ingest of a large number of files is needed by a customer environment.

- File access latencies are affected by the number of files in a directory. To the extent possible, we recommend directory sizes of less than 250,000. Larger directory sizes might experience slower responses to metadata operations such as listing the files in the directory and opening or creating a file.

Best practices for data streams sent to DD systems

For optimal performance, Dell recommends limits on simultaneous streams between DD systems, and your backup servers.

A data stream, in the context of the following table, refers to a large byte stream associated with sequential file access, such as a write stream to a backup file or a read stream from a restore image. A Replication source or destination stream refers to a replication operation or a DD Boost file replication stream associated with a file replication operation.

Table 8. Data streams sent to a protection system

Model	RAM / NVRAM	Backup write streams	Backup read streams	Repl[a] source streams	Repl[a] dest streams	Mixed
DD6300	48 or 96 GB / 8 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest+w<=270; Total<=270
DD6400	192 GB/ 16 GB	270	75	150	270	w<=270; r<=75; ReplSrc<=150; ReplDest<=270; ReplDest+w<=270; Total<=270
DD6800	192 GB / 8 GB	400	110	220	400	w<=400; r<=110; ReplSrc<=220; ReplDest<=400; ReplDest+w<=400; Total<=400
DD6900	288 GB / 16 GB	400	110	220	400	w<=400; r<=110; ReplSrc<=220; ReplDest<=400; ReplDest+w<=400; Total<=400
DD9300	192 or 384 GB / 8 GB	800	220	440	800	w<=800; r<=220; ReplSrc<=440; ReplDest<=800; ReplDest+w<=800; Total<=800
DD9400	576 GB / 16 GB	800	220	440	800	w<=800; r<=220; ReplSrc<=440; ReplDest<=800; ReplDest+w<=800; Total<=800
DD9800	256 or 768 GB / 8 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest+w<=1080; Total<=1885
DD9900	1152 GB / 16 GB	1885	300	540	1080	w<=1885; r<=300; ReplSrc<=540; ReplDest<=1080; ReplDest+w<=1080; Total<=1885
DD VE 8 TB	8 GB / 512 MB	20	16	20	20	w<= 20 ; r<= 16 ReplSrc<=20; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=20;Total<=20
DD VE 16 TB	16 GB / 512 MB or 24 GB / 1 GB	45	30	45	45	w<= 45 ; r<= 30 ReplSrc<=45; ReplDest<=45; ReplDest+w<=45; w+r+ReplSrc <=45;Total<=45
DD VE 32 TB	24 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 48 TB	36 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90

Table 8. Data streams sent to a protection system (continued)

Model	RAM / NVRAM	Backup write streams	Backup read streams	Repl[a] source streams	Repl[a] dest streams	Mixed
DD VE 64 TB	48 GB / 1 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD VE 96 TB	64 GB / 2 GB	180	50	90	180	w<= 180 ; r<= 50 ReplSrc<=90; ReplDest<=180; ReplDest+w<=180; w+r+ReplSrc <=180;Total<=180
DD3300 4 TB	12 GB (virtual memory) / 512 MB	20	16	30	20	w<= 20 ; r<= 16 ReplSrc<=30; ReplDest<=20; ReplDest+w<=20; w+r+ReplSrc <=30;Total<=30
DD3300 8 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 16 TB	32 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=90
DD3300 32 TB	46 GB (virtual memory) / 1.536 GB	90	50	90	90	w<= 90 ; r<= 50 ReplSrc<=90; ReplDest<=90; ReplDest+w<=90; w+r+ReplSrc <=90;Total<=140

a. DirRepl, OptDup, MTreeRepl streams

Monitoring the file system

View real-time data storage statistics and manage file system cleaning, expansion, copying, and destruction.

About this task

The File System view provides the following tabs for viewing details:

- Status—Display the status of file system services.
- Summary—Shows space usage statistics for the active and cloud tiers and enables you to view file system status, configure file system settings, and perform Fast Copy, expand capacity, and destroy file system operations. For Cloud Tier, the **Cloud File Recall** field contains a **Recall** link to initiate a file recall from the Cloud Tier. A **Details** link is available if any active recalls are underway. For more information, see the "Recalling a File from the Cloud Tier" topic.
- Settings—Display and change system options as well as the current cleaning schedule.
- Cloud Units—Display summary information for cloud units, add and modify cloud units, and manage certificates. Shown only when the optional Cloud Tier license is enabled. This view lists summary information (status, data movement throttle, read access, local compression, data movement and data status) the name of the cloud provider, the used capacity, and the licensed capacity. Controls are provided for editing the cloud unit, managing certificates, and adding a new cloud unit.
- Encryption—Display encryption status, progress, algorithms, and so on.
- Space Usage—Display a visual (but static) representation of data use for the file system at certain points in time.
- Consumption—Display space used over time, in relation to total system capacity.
- Daily Written—Display the flow of data over time. The data amounts are shown over time for pre- and post-compression amounts.

Steps

- Select **Data Management** > **File System**.

Managing file system capacity

DD systems have three progressive levels of capacity. As each level is reached, more operations are progressively disallowed. At each level, deleting data and then performing a file system cleaning operation makes disk space available.

NOTE: The process of deleting files and removing snapshots does not immediately reclaim disk space, but the next cleaning operation reclaims the space.

- Level 1—At the first capacity level, no more new data can be written to the file system. An informative out-of-space alert is generated.

  Remedy—Delete unneeded data sets, reduce the retention period, delete snapshots, and perform a file system cleaning operation.

- Level 2—At the second capacity level, files cannot be deleted because deleting files also requires free space. At this level, the system does not have enough free space available to delete files.

  Remedy—Expire snapshots and perform a file system cleaning operation.

- Level 3—At the third and final capacity level, attempts to expire snapshots, delete files, or write new data fail.

  Remedy—Perform a file system cleaning operation to free enough space to delete some files or expire some snapshots and then rerun cleaning.

Monitor the capacity with email alerts

Alerts are generated when the file system is at 75% (for `/ddvar`), 90%, 95%, and 100% full. To send these alerts to a specific user, add the user to the alert emailing list.

See Managing Alerts.

Related concepts

Health Alerts panel

Performing basic operations

Basic file system operations include enabling and disabling the file system, and in the rare occasion, destroying a file system.

Creating the file system

Create a file system from the Data Management > File System page using the Summary tab.

About this task

There are three reasons to create a file system:

- For a new system.
- When a system is started after a clean installation.
- After a file system has been destroyed.

To create the file system:

Steps

1. Verify that storage has been installed and configured (see the section on viewing system storage information for more information). If the system does not meet this prerequisite, a warning message is displayed. Install and configure the storage before attempting to create the file system.
2. Select **Data Management** > **File System** > **Summary** > **Create**.

   The File System Create Wizard is launched. Follow the instructions provided.

Related tasks

Viewing system storage information

Enabling or disabling the file system

The option to enable or disable the file system is dependent on the current state of the file system—if its enabled, you can disable it and vice versa.

About this task

- Enabling the file system allows system operations to begin. This ability is available to administrative users only.
- Disabling the file system halts all system operations, including cleaning. This ability is available to administrative users only.

CAUTION: Disabling the file system when a backup application is sending data to the system can cause the backup process to fail. Some backup software applications are able to recover by restarting where they left off when they are able to successfully resume copying files; others might fail, leaving the user with an incomplete backup.

Steps

1. Select **Data Management** > **File System** > **Summary**.
2. For **File System**, click **Enable** or **Disable**.
3. On the confirmation dialog, click **Close**.

Expanding the file system

You might need to expand the size of a file system if the suggestions given in "When the File System Is Full or Nearly Full" do not clear enough space for normal operations.

About this task

A file system may not be expandable, however, for these reasons:

- The file system is not enabled.
- There are no unused disks or enclosures in the Active or Cloud tiers.
- An expanded storage license is not installed.
- There are not enough capacity licenses installed.

DD6300 systems support the option to use ES30 enclosures with 4 TB drives (43.6 TiB) at 50% utilization (21.8 TiB) in the active tier if the available licensed capacity is exactly 21.8 TiB. The following guidelines apply to using partial capacity shelves.

- No other enclosure types or drive sizes are supported for use at partial capacity.
- A partial shelf can only exist in the Active tier.
- Only one partial ES30 can exist in the Active tier.
- Once a partial shelf exists in a tier, no additional ES30s can be configured in that tier until the partial shelf is added at full capacity.

  NOTE: This requires licensing enough additional capacity to use the remaining 21.8 TiB of the partial shelf.
- If the available capacity exceeds 21.8 TB, a partial shelf cannot be added.
- Deleting a 21 TiB license will not automatically convert a fully-used shelf to a partial shelf. The shelf must be removed, and added back as a partial shelf.

For DD6900, DD9400, and DD9900 systems, storage capacity licenses are available in increments of 60 TB raw (48 TB usable) capacity. Therefore, systems with 8 TB drives may encounter situations where the licensed capacity does not match the full capacity of the disks installed in the disk shelves. For example, if a system has a licensed capacity of 48 TB usable capacity, and has one pack of 8 TB disks for a total of 96 TB usable capacity, only half the system capacity is available for use.

For DD9400 and DD9900 systems, mixing 4 TB and 8 TB drives may not allow the system to reach its maximum supported capacity.

To expand the file system:

Steps

1. Select **Data Management** > **File System** > **Summary** > **Expand Capacity**.

   The Expand File System Capacity wizard is launched. The **Storage Tier** drop-down list always contains Active Tier, and it may contain Cloud Tier as a secondary choice. The wizard displays the current capacity of the file system for each tier as well as how much additional storage space is available for expansion.

   NOTE: File system capacity can be expanded only if the physical disks are installed on the system and file system is enabled.

2. From the **Storage Tier** drop-down list, select a tier.
3. In the **Addable Storage** area, select the storage devices to use and click **Add to Tier**.
4. Follow the instructions in the wizard. When the confirmation page is displayed, click **Close**.

Related concepts

Managing file system capacity

Destroying the file system

Destroying the file system should be done only after careful consideration of the ramifications. This action deletes all data in the file system, including virtual tapes. Deleted data is not recoverable. This operation also removes Replication configuration settings.

About this task

This operation is used when it is necessary to clean out existing data, to create a new collection replication destination, or to replace a collection source, or for security reasons because the system is being removed from operation.

NOTE: As this is a destructive procedure, this operation is available to administrative users only and requires security officer authorization. The operation is not allowed if a security policy is not configured. If multifactor authentication is enabled, the security officer must enter an RSA SecurID token after their password.

Steps

Run the `filesys destroy` command.

Performing cleaning

This section provides information about cleaning and describes how to start, stop, and modify cleaning schedules.

DDOS attempts to maintain a counter called 'Cleanable GiB' for the active tier. This number is an estimation of how much physical (postcomp) space could potentially be reclaimed in the active tier by running clean/garbage collection. This counter is shown using the `filesys show space` and `df` commands.

```
Active Tier:
Resource Size GiB Used GiB Avail GiB Use% Cleanable GiB*
---------------- -------- --------- --------- ----
--------------
/data: pre-comp - 7259347.5 - - -
/data: post-comp 304690.8 251252.4 53438.5 82% 51616.1 <=== NOTE
/ddvar 29.5 12.5 15.6 44% -
---------------- -------- --------- --------- ----
--------------
```

Run active tier clean if either:

- The value for 'Cleanable GiB' is large
- DDFS has become 100% full (and is therefore read-only)

Clean may not reclaim all potential space in a single run. On systems containing very large datasets, clean works against the portion of the file system containing the most superfluous data and may need to be run multiple times before all potential space is reclaimed.

Capacity prediction-enabled automatic cleaning

Prediction-enabled automatic cleaning complements the existing cleaning engine by predicting system capacity and allowing cleaning to start automatically when the system predicts it will meet certain capacity usage levels instead of relying entirely on time-based cleaning schedules regardless of capacity usage or system activity.

On systems that do not have high capacity usage, running a weekly scheduled cleaning can impact system performance for minimal space reclamation. This feature predicts the plausible capacity usage of the system over time based on the current capacity usage and data ingest trends. Automatic cleaning uses this prediction engine to start cleaning on the Active tier when it detects the system will exceed the specified automatic cleaning thresholds.

Automatic cleaning collects and collates different data parameters at a defined interval. If sufficient data points are not available to efficiently predict the capacity usage, the system defaults to the active time-based cleaning schedule with no user intervention required, even when the auto-clean feature is enabled. In a deduplication system, capacity usage does not follow a straight line or grow linearly, making capacity prediction more difficult. The available data on the system is the difference of total data ingested and total data cleaned. The total ingest may follow a linear or partially linear progression, while the amount of data cleaned may be more arbitrary as users can delete different amounts of data. Instead of predicting total system capacity against time, automatic cleaning models a sophisticated algorithm that provides a prediction on when the system will reach the capacity margin specified by the automatic cleaning policy. If the time is more than the number of days set in the policy, cleaning will not run. If the data ingest rate changes, the algorithm adapts to the new data and adjusts the prediction.

Prediction-based cleaning must be manually enabled. There are two modes of operation:

- Fully automatic cleaning
- Scheduled automatic cleaning

Fully automatic cleaning

Fully automatic cleaning is only supported on systems that do not have Cloud Tier configured. Run the following command to configure fully automatic cleaning:

filesys clean auto schedule estimate-percent-used *<capacity-percentage>* days *<number-of-days-in-estimation-period>* interval-days *<number-of-days>*

```
# filesys clean auto schedule estimate-percent-used 92 days 14 interval-days 7

This will start scheduled cleaning when prediction engine estimates post-comp used
capacity to reach 92%, in next 14 days - 7 days is the  interval between two consecutive
GC runs.
```

Scheduled automatic cleaning

Scheduled automatic cleaning is required for systems with Cloud Tier storage. Every cleaning run that is skipped on the Active Tier is considered as a successful run to allow cloud cleaning operations to proceed on their assigned schedules.

At the scheduled time, the system queries the prediction engine to check the projected estimation for capacity usage and starts cleaning on the Active Tier if capacity usage is projected to meet or exceed the specified threshold. If the capacity usage is projected to be under the specified threshold, the cleaning operation is skipped.

Run the following command to configure scheduled automatic cleaning:

filesys clean skip schedule estimate-percent-used *<capacity-percentage>* days *<number-of-days-in-estimation-period>*

```
# filesys clean skip schedule estimate-percent-used 92 days 14

As per the set schedule(filesys clean set schedule), on the scheduled time & day cleaning
starts if the prediction engine estimates post-comp used capacity to reach 92%, in next 14
days. The interval is as per the set schedule.
```

Starting cleaning

To immediately start a cleaning operation.

About this task

If the system detects an anomaly in the amount of deleted data, scheduled cleaning operations will be skipped. Manual cleaning operations are available, but will require a security policy to exist on the system, and security officer authorization when the command to initiate cleaning is run.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Settings** > **Cleaning**.
   The **Cleaning** tab of the **File System Setting** dialog displays the configurable settings for each tier.
2. For the Active Tier:
   a. In the **Throttle %** field, enter a system throttle amount. This is the percentage of CPU usage dedicated to cleaning. The default is 50 percent.
   b. In the **Frequency** drop-down list, select one of these frequencies: **Never**, **Daily**, **Weekly**, **Biweekly**, or **Monthly**. The default is **Weekly**.

      CAUTION: Selecting Daily may impact future replication and restore operations. Dell recommends running cleaning operations no more than twice per week.

   c. For **At**, configure a specific time.
3. For the Cloud Tier:
   a. In the Throttle % text box, enter a system throttle amount. This is the percentage of CPU usage dedicated to cleaning. The default is 50 percent.
   b. In the Frequency drop-down list, select one of these frequencies: Never, After every 'N' Active Tier cleans.

      NOTE: If a cloud unit is inaccessible when cloud tier cleaning runs, the cloud unit is skipped in that run. Cleaning on that cloud unit occurs in the next run if the cloud unit becomes available. The cleaning schedule determines the duration between two runs. If the cloud unit becomes available and you cannot wait for the next scheduled run, you can start cleaning manually.

4. Click **Save**.

   NOTE:

   To start the cleaning operation using the CLI, use the `filesys clean start` command.

   ```
   # filesys clean start
   Cleaning started. Use 'filesys clean watch' to monitor progress.
   ```

   To confirm that cleaning is in progress, use the `filesys status` command.

   ```
   # filesys status
   The filesystem is enabled and running.
   Cleaning started at 2017/05/19 18:05:58: phase 1 of 12 (pre-merge)
   50.6% complete, 64942 GiB free; time: phase 0:01:05, total 0:01:05
   ```

   If cleaning is already running, the following message is displayed when it is attempted to be started.

   ```
   **** Cleaning already in progress. Use 'filesys clean watch' to monitor progress.
   ```

   NOTE: If clean is not able to start, contact the contracted support provider for further assistance. This issue may indicate that the system has encountered a `missing segment error`, causing clean to be disabled.

Scheduling or stopping cleaning

To immediately stop or schedule a cleaning operation.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Settings** > **Cleaning**.
   The Cleaning tab of the File System Setting dialog displays the configurable settings for each tier.
2. For the active tier:
   a. In the Frequency drop-down list, select wanted frequency.
3. For the cloud tier:
   a. In the Frequency drop-down list, select wanted frequency.
4. Click **Save**.

NOTE: The CLI can be used to check that a clean schedule has been set.

```
# filesys clean show schedule
```

If necessary, set an active tier clean schedule. The following example sets cleaning to run every Tuesday at 6 AM:

```
# filesys clean set schedule Tue 0600
Filesystem cleaning is scheduled to run "Tue" at "0600".
```

Performing sanitization

To comply with government guidelines, system sanitization, also called data shredding, must be performed when classified or sensitive data is written to any system that is not approved to store such data.

When an incident occurs, the system administrator must take immediate action to thoroughly eradicate the data that was accidentally written. The goal is to effectively restore the storage device to a state as if the event never occurred. If the data leakage is with sensitive data, the entire storage will need to be sanitized using Dell Professional Services' Secure Data erasure practice.

The sanitization command exists to enable the administrator to delete files at the logical level, whether a backup set or individual files. Deleting a file in most file systems consists of just flagging the file or deleting references to the data on disk, freeing up the physical space to be consumed at a later time. However, this simple action introduces the problem of leaving behind a residual representation of underlying data physically on disks. Deduplicated storage environments are not immune to this problem.

System sanitization requires security officer authorization, and cannot be run if a security policy is not configured on the system. KB article 545871 *PowerProtect Data Domain Operating Systems - Added Protection Against Accidental Execution of Commands*, available from https://www.dell.com/support, provides more information.

Shredding data in a system implies eliminating the residual representation of that data and thus the possibility that the file may be accessible after it has been shredded. Dell's sanitization approach ensures it is compliant with the *National Institute of Systems and Technology (NIST) Special Publication 800-88 Guidelines for Media Sanitization*.

Sanitizing deduplicated data

Protection systems sanitize data in place, in its native deduplicated state.

Deduplication storage systems extract common data patterns from files sent to the system and store only unique copies of these patterns, referencing all the redundant instances. Because these data patterns or segments may potentially be shared among many files in the system, the sanitization process must first determine whether each of the segments of the contaminated file are shared with a clean file and then erase only those segments that are not shared, along with any contaminated metadata.

All storage tiers, caches, unused capacity, and free space are cleared so that every copy of every segment that belongs exclusively to the deleted files is eradicated. The system reclaims and overwrites all of the storage occupied by these segments to effectively restore the storage device to a state as if the contaminated files never existed in that system.

Sanitization level 1: data clearing or shredding

If the data you need to remove is unclassified, as defined in the "US Department of Defense 5220.22-M Clearing and Sanitization Matrix," Level 1 sanitization can be used to overwrite the affected storage once. This provides the basis for handling most data shredding and system sanitization cases.

About this task

The system sanitization feature ensures that every copy of every segment that belongs only to erased files is overwritten using a single-pass zerotization mechanism. Clean data in the system being sanitized is online and available to users.

Steps

1. Delete the contaminated files or backups through the backup software or corresponding client. In the case of backups, be sure to manage the backup software appropriately to ensure that related files on that image are reconciled, catalog records are managed as required, and so forth.
2. Run the `system sanitize start` command on the contaminated system to cause all previously used space in it to be overwritten once. This command requires security officer authorization to run, and cannot run if a security policy is not configured on the system. If multifactor authentication is enabled, the security officer must enter an RSA SecurID token after their password.

   NOTE: The `system sanitize start` command does not run on Cloud Tier-enabled systems.

3. Wait for the affected system to be sanitized. Sanitization can be monitored by using the `system sanitize watch` command.

   If the affected system has replication enabled, all the systems containing replicas need to be processed in a similar manner. Depending on how much data exists in the system and how it is distributed, the `system sanitize` command could take some time. However, during this time, all clean data in the system is available to users.

Sanitization level 2: full system sanitization

If the data you need to remove is classified, as defined in the *NIST Special Publication 800-88 Guidelines for Media Sanitization*, level 2 sanitization (full system sanitization) is now required.

About this task

Contact Dell Support for additional information about level 2 system sanitization.

Modifying basic settings

Change the type of compression used, marker types, Replica write status, and Staging Reserve percentage, as described in this section.

Changing local compression

Use the General tab of the File System Settings dialog to configure the local compression type.

About this task

NOTE: Do not change the type of local compression unless it is necessary.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Settings** > **General**.
2. From the Local Compression Type drop-down list, select a compression type.

Table 9. Compression type

Option	Description
NONE	Do not compress data.
LZ	The algorithm that gives the best throughput. Dell recommends the lz option, which is the default setting, for the following systems: ● DD3300 ● DD6300 ● DD6800 ● DD9300 ● DD9800
GZFAST	A zip-style compression that uses less space for compressed data, but more CPU cycles (twice as much as lz). Gzfast is the recommended alternative for sites that want more compression at the cost of lower performance. Dell recommends the gzfast option, which is the default setting, for the following systems: ● DD6400 ● DD6900 ● DD9400 ● DD9900
GZ	A zip-style compression that uses the least amount of space for data storage (10% to 20% less than lz on average; however, some datasets get much higher compression). This also uses the most CPU cycles (up to five times as much as lz). The gz compression type is commonly used for nearline storage applications in which performance requirements are low.

3. Click **Save**.

Types of compression

DDOS compresses data at two levels: global and local. Global compression compares received data to data already stored on disks. Duplicate data does not need to be stored again, while data that is new is locally compressed before being written to disk.

Local Compression

A protection system uses a local compression algorithm developed specifically to maximize throughput as data is written to disk. The lz algorithm allows shorter backup windows for backup jobs but uses more space. Two other types of local compression are available, gzfast and gz. Both provide increased compression over lz, but at the cost of additional CPU load. Local compression options provide a trade-off between slower performance and space usage. It is also possible to turn off local compression. To change compression, see Changing local compression.

After you change the compression, all new writes use the new compression type. Existing data is converted to the new compression type during cleaning. It takes several rounds of cleaning to recompress all of the data that existed before the compression change.

The initial cleaning after the compression change might take longer than usual. Whenever you change the compression type, carefully monitor the system for a week or two to verify that it is working properly.

DD3300, DD6300, DD6800, DD9300, and DD9800 systems use the lz compression algorithm as the default local compression type.

DD6900, DD9400, and DD9900 systems use the gzfast algorithm as the default local compression type.

Intel Quick Assist Technology (QAT)

PowerProtect DD6900, DD9400, and DD9900 systems come with a hardware accelerator card to support Intel QAT in combination with gzfast compression. DDOS offloads compression and decompression work to the hardware accelerator to achieve a higher compression ratio and free up CPU resources to improve system performance. No additional configuration is required.

If the QAT card encounters a problem and cannot function normally, it fails over to the CPU. The system generates an alert if the QAT fails over more than 10,000 time per hour for 8 hours.

The QAT also has a memory dump functionality that saves patterns that lead to a QAT hardware hang or failure into a log file for future analysis. This log file is located in `/ddr/var/core/qat_dump_file/`. After a QAT memory dump, the completed dump log is saved to `/ddr/var/core/qat_mem_dump.log`.

Changing read-only settings

Change the replica to writable. Some backup applications must see the replica as writable to do a restore or vault operation from the replica.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Settings** > **General**.
2. In the Report Replica as Writable area, toggle between **Disabled** and **Enabled** as appropriate.
3. Click **Save**.

Working with disk staging

Disk staging enables a protection system to serve as a staging device, where the system is viewed as a basic disk via a CIFS share or NFS mount point.

Disk staging can be used in conjunction with your backup software, such as NetWorker and Veritas NetBackup (NBU), it does not require a license, and is disabled by default.

NOTE: The DD VTL feature is not required or supported when the system is used as a Disk Staging device.

The reason that some backup applications use disk staging devices is to enable tape drives to stream continuously. After the data is copied to tape, it is retained on disk for as long as space is available. Should a restore be needed from a recent backup, more than likely the data is still on disk and can be restored from it more conveniently than from tape. When the disk fills up, old backups can be deleted to make space. This delete-on-demand policy maximizes the use of the disk.

In normal operation, the system does not reclaim space from deleted files until a cleaning operation is done. This is not compatible with backup software that operates in a staging mode, which expects space to be reclaimed when files are deleted. When you configure disk staging, you reserve a percentage of the total space—typically 20 to 30 percent—in order to allow the system to simulate the immediate freeing of space.

The amount of available space is reduced by the amount of the staging reserve. When the amount of data stored uses all of the available space, the system is full. However, whenever a file is deleted, the system estimates the amount of space that will be recovered by cleaning and borrows from the staging reserve to increase the available space by that amount. When a cleaning operation runs, the space is actually recovered and the reserve restored to its initial size. Since the amount of space made available by deleting files is only an estimate, the actual space reclaimed by cleaning may not match the estimate. The goal of disk staging is to configure enough reserve so that you do not run out before cleaning is scheduled to run.

Configuring disk staging

Enable disk staging and specify the staging reserve percentage.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Settings** > **General**.
2. In the Staging Reserve area, toggle between **Disabled** and **Enabled** as appropriate.
3. If Staging Reserve is enabled, enter a value in the % of Total Space box.
   This value represents the percentage of the total disk space to be reserved for disk staging, typically 20 to 30%.
4. Click **Save**.

Tape marker settings

Backup software from some vendors insert markers (tape markers, tag headers, or other names are used) in all data streams (both file system and DD VTL backups) sent to a protection system.

Markers can significantly degrade data compression. As such, the default marker type auto is set and cannot be changed by the user. If this setting is not compatible with your backup software, contact your contracted support provider.

SSD Random workload share

The value for the threshold at which to cap random I/O on the protection system can be adjusted from the default value to accommodate changing requirements and I/O patterns.

By default, the SSD random workload share is set at 40%. This value can be adjusted up or down as needed. Select **Data Management** > **File System** > **Summary** > **Settings** > **Workload Balance**, and adjust the slider.

Click **Save**.

Fast copy operations

A fast copy operation clones files and directory trees of a source directory to a target directory on a protection system.

The `force` option allows the destination directory to be overwritten if it exists. Executing the fast copy operation displays a progress status dialog box.

NOTE: A fast copy operation makes the destination equal to the source, but not at a specific time. There are no guarantees that the two are or were ever equal if you change either folder during this operation.

Performing a fast copy operation

Copy a file or directory tree from a protection system source directory to another destination on the same system.

Steps

1. Select **Data Management** > **File System** > **Summary** > **Fast Copy**.
   The Fast Copy dialog is displayed.
2. In the Source text box, enter the pathname of the directory where the data to be copied resides. For example, `/data/col1/backup/.snapshot/`*snapshot-name*`/dir1`.

   NOTE: col1 uses a lower case L followed by the number 1.

3. In the Destination text box, enter the pathname of the directory where the data will be copied to. For example, `/data/col1/backup/dir2`. This destination directory must be empty, or the operation fails.
   - If the Destination directory exists, click the checkbox **Overwrite existing destination if it exists**.

   CAUTION: This option deletes all data in the destination directory.

4. Click **OK**.
5. In the progress dialog box that appears, click **Close** to exit.

14

Managing MTrees

This chapter includes:

Topics:

MTrees overview

An MTree is a logical partition of the file system.

You can use MTrees for CIFS shares, DD Boost storage units, DD VTL pools, or NFS exports. MTrees allow granular management of snapshots, quotas, and DD Retention Lock.

NOTE:

The maximum number of configurable MTrees on the system can be designated for MTree replication contexts.

Do not place user files in the top-level directory of an MTree. Create subdirectories within the MTree to store user data.

MTree limits

MTree limits for DD systems

Table 10. Number of supported MTrees

System	DDOS Version	Supported configurable MTrees	Supported concurrently active MTrees
DD9900	7.0 and later	256	256
DD6900, DD9400	7.0 and later	128	128
DD6400	7.7 and later	128	128
DD9800	6.0 and later	256	256
DD6800, DD9300	6.0 and later	128	128
DD6300	6.0 and later	100	32

MTree Quotas

MTree quotas apply only to the logical data written to the MTree.

An administrator can set storage space restrictions for an MTree, Storage Unit, or DD VTL pool to prevent it from consuming excess space. MTrees have hard quota limits and soft quota limits. You can set soft, hard, or a combination of both limits. The values must be integers, and the soft value must be less than the hard value.

When a soft limit is set, an alert is generated when the MTree size exceeds the limit, but data can still be written to it. When a hard limit is set and the limit is reached, data cannot be written to the MTree and all write operations fail until data is deleted from the MTree.

See Configure MTree quotas for more information.

Quota enforcement

Enable or disable quota enforcement.

Viewing and monitoring MTree usage

You can view active MTrees and real-time data storage statistics. Information in the overview area is helpful for visualizing space usage trends.

Select **Data Management** > **MTree**.

The MTree view shows a list of configured MTrees. Details of the selected MTree are shown in the **Summary** tab. The **Space Usage** and **Daily Written** tabs show graphs that visually display space usage amounts and data written trends for the selected MTree. The view also contains options that enable MTree configuration for CIFS, NFS, and DD Boost, as well as sections for managing snapshots and DD Retention Lock for an MTree.

To filter for specific MTree names, enter text (wildcards are supported) in the **Filter By MTree Name** box and click **Update**.

NOTE: Physical capacity measurement (PCM) provides space usage information for MTrees. For more information about PCM, see the section regarding understanding physical capacity measurement.

Monitoring MTree usage

Display space usage and data written trends for an MTree.

Steps

- Select **Data Management** > **MTree**.

  The MTree view shows a list of configured MTrees. Details of the selected MTree are shown in the **Summary** tab. The **Space Usage** and **Daily Written** tabs show graphs that visually display space usage amounts and data written trends for the selected MTree. The view also contains options that enable MTree configuration for CIFS, NFS, and DD Boost, as well as sections for managing snapshots and DD Retention Lock for an MTree.

  NOTE: Physical capacity measurement (PCM) provides space usage information for MTrees. For more information about PCM, see the section regarding understanding physical capacity measurement.

Understanding physical capacity measurement

Physical capacity measurement (PCM) provides space usage information for a sub-set of storage space.

From the DD System Manager, PCM provides space usage information for MTrees. Using the command line interface you can view space usage information for MTrees, tenants, tenant units, and path sets.

Once a path is selected for PCM, all paths under it are automatically included.

The *DDOS Command Reference Guide* provides more information about using PCM from the command line.

Enabling, disabling, and viewing physical capacity measurement

Physical capacity measurement provides space usage information for an MTree.

Steps

1. Select **Data Management** > **File System** > **Summary**.

   The system displays the Summary tab in the File System panel.
2. Click **^** in the bottom-right corner to view the status panel.
3. Click **Enable** to the right of **Physical Capacity Measurement Status** to enable PCM.
4. Click **Details** to the right of **Physical Capacity Measurement Status** to view currently running PCM jobs.

- **MTree**: The MTree that PCM is measuring.
- **Priority**: The priority (normal or urgent) for the task.
- **Submit Time**: The time the task was requested.
- **Duration**: The length of time PCM ran to accomplish of the task.

5. Click **Disable** to the right of **Physical Capacity Measurement Status** to disable PCM and cancel all currently running PCM jobs.

Related tasks

Initializing physical capacity measurement
Starting physical capacity measurement immediately

Initializing physical capacity measurement

Physical capacity measurement (PCM) initialization is a one-time action that can take place only if PCM is enabled and the cache has not been initialized. It cleans the caches and enhances measuring speed. During the initialization process, you can still manage and run PCM jobs.

About this task

PCM will not run if the system capacity utilization is more than 90%.

Steps

1. Select **Data Management** > **File System** > **Configuration**.
2. Click **Initialize** under Physical Capacity Measurement to the right of Cache.
3. Click **Yes**.

Managing physical capacity measurement schedules

Create, edit, delete, and view physical capacity measurement schedules. This dialog only displays schedules created for MTrees and schedules that currently have no assignments.

Steps

1. Select **Data Management** > **MTree** > **Manage Schedules**.
   - Click **Add** (+) to create a schedule.
   - Select a schedule and click **Modify** (pencil) to edit the schedule.
   - Select a schedule and click **Delete** (X) to delete a schedule.
2. Optionally, click the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Related tasks

Creating physical capacity measurement schedules
Editing physical capacity measurement schedules
Assigning physical capacity measurement schedules to an MTree

Creating physical capacity measurement schedules

Create physical capacity measurement schedules and assign them to MTrees.

Steps

1. Select **Data Management** > **MTree** > **Manage Schedules**.
2. Click **Add** (+) to create a schedule.
3. Enter the name of the schedule.
4. Select the status:

- **Normal**: Submits a measurement task to the processing queue.
- **Urgent**: Submits a measurement task to the front of the processing queue.

5. Select how often the schedule triggers a measurement occurrence: every **Day**, **Week**, or **Month**.
   - For **Day**, select the time.
   - For **Week**, select the time and day of the week.
   - For **Month**, select the time, and days during the month.
6. Select MTree assignments for the schedule (the MTrees that the schedule will apply to):
7. Click **Create**.
8. Optionally, click on the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Related tasks

Editing physical capacity measurement schedules
Assigning physical capacity measurement schedules to an MTree

Editing physical capacity measurement schedules

Edit a physical capacity measurement schedule.

Steps

1. Select **Data Management** > **MTree** > **Manage Schedules**.
2. Select a schedule and click **Modify** (pencil).
3. Modify the schedule and click **Save**.
   Schedule options are described in the Creating physical capacity measurement schedules topic.
4. Optionally, click the heading names to sort by schedule: **Name**, **Status** (Enabled or Disabled) **Priority** (Urgent or Normal), **Schedule** (schedule timing), and **MTree Assignments** (the number of MTrees the schedule is assigned to).

Related tasks

Assigning physical capacity measurement schedules to an MTree

Assigning physical capacity measurement schedules to an MTree

Attach schedules to an MTree.

Prerequisites

Physical capacity measurement (PCM) schedules must be created.

About this task

NOTE: Administrators can assign up to three PCM schedules to an MTree.

Steps

1. Select **Data Management** > **MTree** > **Summary**.
2. Select MTrees to assign schedules to.
3. Scroll down to the Physical Capacity Measurements area and click **Assign** to the right of Schedules.
4. Select schedules to assign to the MTree and click **Assign**.

Related tasks

Starting physical capacity measurement immediately

Starting physical capacity measurement immediately

Start the measurement process as soon as possible.

Steps

1. Select **Data Management** > **MTree** > **Summary**.
2. Scroll down to the Physical Capacity Measurements area and click **Measure Now** to the right of Submitted Measurements.
3. Select **Normal** (Submits a measurement task to the processing queue), or **Urgent** (Submits a measurement task to the front of the processing queue).
4. Click **Submit**.

Setting the physical capacity measurement throttle

Set the percentage of system resources that are dedicated to physical capacity measurement.

Steps

1. Select **Data Management** > **File System** > **Settings**.
2. In the Physical Capacity Measurement area, click **Edit** to the left of Throttle.
3.

Option	Description
Click Default	Enters the 20% system default.
Type throttle percent	The percentage of system resources that are dedicated to physical capacity measurement.

4. Click **Save**.

Managing MTree operations

This section describes MTree creation, configuration, how to enable and disable MTree quotas, and so on.

Creating an MTree

An MTree is a logical partition of the file system. Use MTrees CIFS shares, DD Boost storage units, DD VTL pools, or NFS exports.

About this task

MTrees are created in the area /data/col1/*mtree_name*.

Steps

1. Select **Data Management** > **MTree**.
2. In the MTree overview area, click **Create**.
3. Enter the name of the MTree in the MTree Name text box. MTree names can be up to 50 characters. The following characters are acceptable:
   - Upper- and lower-case alphabetical characters: A-Z, a-z
   - Numbers: 0-9
   - Embedded space
   - comma (,)
   - period (.), as long as it does not precede the name.
   - explanation mark (!)
   - number sign (#)
   - dollar sign ($)
   - per cent sign (%)

- plus sign (+)
- at sign (@)
- equal sign (=)
- ampersand (&)
- semi-colon (;)
- parenthesis [(and)]
- square brackets ([and])
- curly brackets ({and})
- caret (^)
- tilde (~)
- apostrophe (unslanted single quotation mark)
- single slanted quotation mark (‘)

4. Set storage space restrictions for the MTree to prevent it from consuming excessive space. Enter a soft or hard limit quota setting, or both. With a soft limit, an alert is sent when the MTree size exceeds the limit, but data can still be written to the MTree. Data cannot be written to the MTree when the hard limit is reached.

   NOTE: The quota limits are pre-compressed values.

   To set quota limits for the MTree, select **Set to Specific value** and enter the value. Select the unit of measurement: MiB, GiB, TiB, or PiB.

   NOTE: When setting both soft and hard limits, a quota’s soft limit cannot exceed the quota’s hard limit.

5. Click **OK**.

   The new MTree displays in the MTree table.

   NOTE: You may need to expand the width of the MTree Name column to see the entire pathname.

Configure and enable/disable MTree quotas

Set the storage space restriction for an MTree, Storage Unit, or DD VTL pool.

The **Data Management** > **Quota** page shows the administrator how many MTrees have no soft or hard quotas set. For MTrees with quotas set, the page shows the percentage of pre-compressed soft and hard limits used.

Consider the following information when managing quotas.

- MTree quotas apply to ingest operations. These quotas can be applied to DD VTL, DD Boost, CIFS, and NFS.
- Snapshots are not counted.
- Quotas cannot be set on the `/data/col1/backup` directory.
- The maximum quota value allowed is 4096 PiB.

Configure MTree quotas

Use the MTree tab or the Quota tab to configure MTree quotas.

About this task

Steps

1. Select one of the following menu paths:
   - Select **Data Management** > **MTree**.
   - Select **Data Management** > **Quota**.
2. Select only one MTree in the MTree tab, or one or more MTrees in the Quota tab.

   NOTE: Quotas cannot be set on the `/data/col1/backup` directory.
3. In the MTree tab, click the **Summary** tab, and then click the **Configure** button in the Quota area.
4. In the Quota tab, click the **Configure Quota** button.

Configuring MTree quotas

Enter values for hard and soft quotas and select the unit of measurement.

Steps

1. In the Configure Quota for MTrees dialog box, enter values for hard and soft quotas and select the unit of measurement: MiB, GiB, TiB, or PiB.
2. Click **OK**.

Deleting an MTree

Removes the MTree from the MTree table. The MTree data is deleted at the next cleaning.

About this task

NOTE: Because the MTree and its associated data are not removed until file cleaning is run, you cannot create a new MTree with the same name as a deleted MTree until the deleted MTree is completely removed from the file system by the cleaning operation.

Steps

1. Select **Data Management** > **MTree**.
2. Select an MTree.
3. In the MTree overview area, click **Delete**.
4. Click **OK** at the Warning dialog box.
5. Click **Close** in the Delete MTree Status dialog box after viewing the progress.

Undeleting an MTree

Undelete retrieves a deleted MTree and its data and places it back in the MTree table.

About this task

An undelete of an MTree retrieves a deleted MTree and its data and places it back in the MTree table.

An undelete is possible only if file cleaning has not been run after the MTree was marked for deletion.

NOTE: You can also use this procedure to undelete a storage unit.

Steps

1. Select **Data Management** > **MTree** > **More Tasks** > **Undelete**.
2. Select the checkboxes of the MTrees you wish to bring back and click **OK**.
3. Click **Close** in the Undelete MTree Status dialog box after viewing the progress.

   The recovered MTree displays in the MTree table.

Renaming an MTree

Use the Data Management MTree GUI to rename MTrees.

Steps

1. Select **Data Management** > **MTree**.
2. Select an MTree in the MTree table.
3. Select the Summary tab.
4. In the Detailed Information overview area, click **Rename**.

5. Enter the name of the MTree in the New MTree Name text box.

   See the section about creating an MTree for a list of allowed characters.

6. Click **OK**.

   The renamed MTree displays in the MTree table.

Related tasks

Creating an MTree

MTree analytics

MTree analytics provides the ability to determine the actual physical space consumed by files within the MTrees on the system after accounting for compression and deduplication.

Due to the deduplication capabilities in DDOS, the total size of a collection of files cannot be calculated by adding the logical sizes of the files because it will count common shared data segments multiple times. This results in a logical size greater than the actual disk space the files occupy. Compression can further reduce disk space usage.

In addition to calculating physical disk space consumption, MTree analytics also provides the ability to determine the unique sizes of the files. This is useful for freeing up space when the file system is full or migrating data to other systems for storage load balancing.

The unique size of a subset of files determines the amount of space available to reclaim when those files are deleted or migrated. For example, if two identical files are 100 GB in size, the unique size of the first file is 0 and deleting it will not reclaim any space in the file system. However, if a 100 GB file does not share data segments with any other file on the system, deleting it will reclaim all 100 GB that it consumes.

After identifying the unique file size, MTree analytics identifies files that are similar and groups them together as a cluster for maximum efficiency. From the example above, if the two identical files were grouped together into a single cluster, deleting, or migrating that cluster would result in 100 GB of space reclamation.

The Global Dedupe Engine (GDE) analyzes and organizes the files on the protection system by sampling all the fingerprints in the file system to determine which files reference each other, computes the MTree statistics, and applies the hierarchical clustering algorithm to create the clusters.

GDE operations consist of five phases:

Merge phase	The system writes all metadata to disk to create a consistency point to prepare for the analysis and enumeration phases. This operation may take up to an hour to complete.
Analysis phase	The system scans all the fingerprints in the file system and constructs a hash vector for the metadata segments. This operation may take a few hours to complete.
Enumeration phase	The metadata segments are scanned and the owner (MTree) of each metadata segment is tracked in the process. At the final step, the live data segment fingerprints, and their owners (MTree) are inserted into a fingerprint to MTree ID dictionary. This procedure can take many hours but usually less than a day. The system scans the metadata segments, tracking the MTree that owns each segment in the process. At the final step, the data segment fingerprints and their MTree ownership information are inserted into a dictionary that maps each fingerprint to its MTree. This operation may take up to a day to complete.
Size estimation phase	The system computes the size of the MTrees. This operation may take up to an hour to complete.
Clustering phase	The system applies the hierarchical clustering algorithm to create clusters of similar MTrees. Optionally, users can specify operating parameters for the GDE. The options are: ● W-index algorithm (default) - Maximize the ratio of the inter-cluster distance divided by the intra-cluster distance. ● Maximum gap - Maximize the difference of the inter-cluster distance minus the intra-cluster distance ● Number of clusters - Stop the algorithm when a specified number of clusters are formed ● Similarity level - Stop the algorithm when the similarity drops below the specified threshold This operation should take a few minutes to complete.

MTree analytics works on all MTrees, regardless of replication configuration.

MTree analytics reports

The results from the MTree analysis are available in two reports.

As MTree analytics runs, it combines the two most similar MTrees into a cluster and repeats the process until only one cluster remains. However, as the process runs, it creates multiple intermediate clusters. The statistics for these clusters are described in the MTree size report.

The MTree size report consists of the following items:

- MTree name or cluster ID - MTree name (for MTrees not combined with any other MTrees in a cluster) or cluster ID (for MTrees that were combined with other MTrees in a cluster

  NOTE: The contents of individual MTree clusters is not displayed.
- Logical size - Total size of the files before local compression and deduplication
- Metadata size - Total physical size of the metadata in the MTree cluster
- Physical size - Total physical of the files in the MTree after deduplication and local compression
- Unique size - Total unique size (physical) of the files in the MTree cluster
- Local compression ratio - Local compression ratio
- Global compression ratio - Total compression ratio, deduplication + local compression
- Percent uniqueness - Unique size / physical size
- Number of files - Number of files

Run the **analytics capacity mtree show mtree-report [Active | *<cloud-unit-name>*]** command to display the MTree size report.

The MTree cluster report displays the optimal cluster solution for the MTrees. It lists the clusters based on the GDE parameter specified at run time.

The first half of the report contains the same information as the MTree size report, with an additional column for the similarity score. - This column displays the intra-cluster distance of the cluster, or the maximum distance among two MTree objects in the cluster.

NOTE: Distance = (1 - similarity).

The second half of the report displays the MTrees in each cluster.

Run the **analytics capacity mtree show cluster-report [Active | *<cloud-unit-name>*] similarity [1-100] ncluster [*<n>*] algo [windex | maxgap]** command to display the cluster report.

Scheduling and running MTree analytics

MTree analytics is set to run after garbage collection on the system. The default schedule for MTree analytics scans the Active Tier is every week and does not scan the Cloud Tier.

About this task

If there are multiple cloud units, MTree analytics alternates between them. For example, if there are two cloud units and the frequency is set to four weeks, the first cloud unit runs on the fourth week and the second cloud unit runs on the eighth week.

MTree Analytics runs after Active Tier garbage collection. It has a low priority to avoid disrupting other GC-related activities including cloud GC, and all verifications. If any of these activities are scheduled to run after GC, the system postpones MTree to the next week. Therefore, it is possible for an analysis to abort and not report any statistics for the aborted cycle.

Run the following commands to use MTree analytics:

Steps

1. Set the MTree analytics schedule.
   a. Run the **analytics capacity mtree frequency set *<n>* [tier active | tier cloud]** command to set the frequency in weeks to run MTree analytics.
   b. Run the **analytics capacity mtree frequency show [tier active | tier cloud]** command to display the current MTree analytics schedule.
2. Run the **analytics capacity mtree start [Active | *<cloud-unit-name>*]** command to run MTree analytics manually.

NOTE: By default, this command starts analysis on the Active Tier.

3. Run the **analytics capacity mtree status [Active | *<cloud-unit-name>*]** command to check the MTree analytics status.

   This command displays:
   - Current, or most recently completed GDE status for the Active Tier or specified cloud unit
   - Analysis start time
   - Analysis end time (if complete)
   - Total time of each GDE phase
   - GDE runtime up to the current phase, or total time if complete
   - Progress (% complete)
   - Phase statistics
4. Run the **analytics capacity mtree watch** command to monitor the progress of an MTree analysis operation in real time.

   NOTE: If no analysis is running, this command does not display an output.
5. Run the **analytics capacity mtree stop** command to stop an in-progress MTree analysis.
6. Run the **analytics capacity mtree frequency reset [tier active | tier cloud]** command to reset the MTree analytics schedule to its default value.

   NOTE: By default, this command resets the Active Tier schedule.

15

Managing Snapshots

This chapter includes:

Topics:

Snapshots overview

A snapshot is a read-only copy of a designated MTree at a specific time. You can use a snapshot as a restore point, you can manage MTree snapshots and schedules, and you can view the status of existing snapshots.

Snapshots for the MTree named `backup` are created in the system directory `/data/col1/backup/.snapshot`. Each directory under `/data/col1/backup` also has a `.snapshot` directory with the name of each snapshot that includes the directory. Each MTree has the same structure. As an example, an MTree named SantaClara would have a system directory `/data/col1/SantaClara/.snapshot`, and each subdirectory in `/data/col1/SantaClara` would have a `.snapshot` directory as well.

Snapshots and CIFS Protocol

As of DDOS 5.0, the `.snapshot` directory is no longer visible in the directory listing in Windows Explorer or DOS CMD shell. You can access the `.snapshot` directory by entering its name in the Windows Explorer address bar or the DOS CMD shell. For example, `\\dd\backup\.snapshot` or `Z:\.snapshot` when `Z:` is mapped as `\\dd\backup`).

Snapshot limitations

- Snapshots created on the source DD system are replicated to the destination system with collection and MTree replication. You cannot create snapshots on a system that is a replica for collection replication. You cannot create a snapshot on the destination MTree of an MTree replication.
- The maximum number of snapshots allowed per MTree is 750. Warnings are sent when the number of snapshots per MTree reaches 90% of the maximum allowed number (from 675 to 749 snapshots), and an alert is generated when the maximum number is reached. To clear the warning, expire snapshots, and then run the file system cleaning operation.

  An expired snapshot remains available until the next file system cleaning operation. To identify an MTree that is nearing the maximum number of snapshots, check the Snapshots panel of the MTree page regarding viewing MTree snapshot information.
- Snapshot retention for an MTree does not take any extra space, but if a snapshot exists and the original file no longer exists, the space cannot be reclaimed.
- The `.snapshot` directory is not visible if only `/data` is mounted. When the MTree itself is mounted, the .snapshot directory is visible.

Related concepts

MTrees overview

Monitoring snapshots and their schedules

This section provides detailed and summary information about the status of snapshots and snapshot schedules.

Managing snapshots

This section describes how to manage snapshots.

Creating a snapshot

Create a snapshot when an unscheduled snapshot is required.

About this task

Steps

1. Select **Data Management** > **Snapshots** to open the Snapshots view.
2. In the Snapshots view, click **Create**.
3. In the Name text field, enter the name of the snapshot.
4. In the MTree(s) area, select a checkbox of one or more MTrees in the Available MTrees panel and click **Add**.
5. In the Expiration area, select one of these expiration options:
   a. **Never Expire**.
   b. Enter a number for the In text field, and select **Days**, **Weeks**, **Month**, or **Years** from the drop-down list. The snapshot will be retained until the same time of day as when it is created.
   c. Enter a date (using the format *mm/dd/yyyy*) in the On text field, or click **Calendar** and click a date. The snapshot will be retained until midnight (00:00, the first minute of the day) of the given date.
6. Click **OK** and **Close**.

Modifying a snapshot expiration date

Modify snapshot expiration dates to remove them or extent their life for auditing or compliance.

Steps

1. Select **Data ManagementSnapshots** to open the Snapshots view.
2. Click the checkbox of the snapshot entry in the list and click **Modify Expiration Date**.

   NOTE: More than one snapshot can be selected by clicking additional checkboxes.

3. In the Expiration area, select one of the following for the expiration date:
   a. **Never Expire**.
   b. In the In text field, enter a number and select **Days**, **Weeks**, **Month**, or **Years** from the drop-down list. The snapshot will be retained until the same time of day as when it is created.
   c. In the **On** text field, enter a date (using the format *mm/dd/yyyy*) or click **Calendar** and click a date. The snapshot will be retained until midnight (00:00, the first minute of the day) of the given date.
4. Click **OK**.

Renaming a snapshot

Use the Snapshot tab to rename a snapshot.

Steps

1. Select **Data Management** > **Snapshots** to open the Snapshots view.

2. Select the checkbox of the snapshot entry in the list and click **Rename**.
3. In the Name text field, enter a new name.
4. Click **OK**.

Expiring a snapshot

Snapshots cannot be deleted. To release disk space, expire snapshots and they will be deleted in the next cleaning cycle after the expiry date.

Steps

1. Select **Data Management** > **Snapshots** to open the Snapshots view.
2. Click the checkbox next to snapshot entry in the list and click **Expire**.

   NOTE: More than one snapshot can be selected by selecting additional checkboxes.

   The snapshot is marked as Expired in the Status column and will be deleted at the next cleaning operation.

Managing snapshot schedules

Set up and manage a series of snapshots that will be automatically taken at regular intervals (a snapshot schedule).

Multiple snapshot schedules can be active at the same time.

NOTE: If multiple snapshots with the same name are scheduled to occur at the same time, only one is retained. Which one is retained is indeterminate, thus only one of the snapshots with that name should be scheduled for a given time.

Creating a snapshot schedule

Create a weekly or monthly snapshot schedule using the Data Management GUI.

Steps

1. Select **Data Management** > **Snapshots** > **Schedules** to open the Schedules view.
2. Click **Create**.
3. In the **Name** text field, enter the name of the schedule.
4. In the **Snapshot Name Pattern** text box, enter a name pattern.

   Enter a string of characters and variables that translates to a snapshot name (for example, **`scheduled-%Y-%m-%d-%H-%m`**, translates to "scheduled-2012-04-12-17-33"). Use alphabetic characters, numbers, _, -, and variables that translate into current values.

5. Click **Validate Pattern & Update Sample**.
6. Click **Next**.
7. Select the date when the schedule will be executed:
   a. Weekly—Click checkboxes next to the days of the week or select **Every Day**.
   b. Monthly—Click the **Selected Days** option and click the dates on the calendar, or select the **Last Day of the Month** option.
   c. Click **Next**.
8. Select the time of day when the schedule will be executed:
   a. At Specific Times—Click **Add** and in the Time dialog that appears, enter the time in the format *hh:mm*, and click **OK**.
   b. In Intervals—Click the drop-down arrows to select the start and end time *hh:mm* and AM or PM. Click the **Interval** drop-down arrows to select a number and then the hours or minutes of the interval.
   c. Click **Next**.
9. In the Retention Period text entry field, enter a number and click the drop-down arrow to select days, months, or years, and click **Next**.

   Schedules must explicitly specify a retention time.

10. Review the parameters in the schedule summary and click **Finish** to complete the schedule or **Back** to change any entries.
11. If an MTree is not associated with the schedule, a warning dialog box asks if you would like to add an MTree to the schedule. Click **OK** to continue (or **Cancel** to exit).
12. To assign an MTree to the schedule, in the MTree area, click the checkbox of one or more MTrees in the Available MTrees panel, then click **Add** and **OK**.

Naming conventions for snapshots created by a schedule

The naming convention for scheduled snapshots is the word scheduled followed by the date when the snapshot is to occur, in the format `scheduled-yyyy-mm-dd-hh-mm`. For example, `scheduled-2009-04-27-13-30`.

The name "mon_thurs" is the name of a snapshot schedule. Snapshots generated by that schedule might have the names `scheduled-2008-03-24-20-00`, `scheduled-2008-03-25-20-00`, etc.

Modifying a snapshot schedule

Change the snapshot schedule name, date, and retention period.

Steps

1. In the schedule list, select the schedule and click **Modify**.
2. In the Name text field, enter the name of the schedule and click **Next**.

   Use alphanumeric characters, and the _ and -.
3. Select the date when the schedule is to be executed:
   a. Weekly—Click checkboxes next to the days of the week or select **Every Day**.
   b. Monthly—Click the **Selected Days** option and click the dates on the calendar, or select the **Last Day of the Month** option.
   c. Click **Next**.
4. Select the time of day when the schedule is to be executed:
   a. At Specific Times—Click the checkbox of the scheduled time in the Times list and click **Edit**. In the Times dialog that appears, enter a new time in the format *hh:mm*, and click **OK**. Or click **Delete** to remove the scheduled time.
   b. In Intervals—Click the drop-down arrows to select the start and end time *hh:mm* and AM or PM. Click the Interval drop-down arrows to select a number and then the hours or minutes of the interval.
   c. Click **Next**.
5. In the Retention Period text entry field, enter a number and click the drop-down arrow to select days, months, or years, and click **Next**.
6. Review the parameters in the schedule summary and click **Finish** to complete the schedule or **Back** to change any entries.

Deleting a snapshot schedule

Delete a snapshot schedule from the schedule list.

Steps

1. In the schedule list, click the checkbox to select the schedule and click **Delete**.
2. In the verification dialog box, click **OK** and then **Close**.

Recover data from a snapshot

Use the fastcopy operation to retrieve data stored in a snapshot. See the section regarding fast copy operations.

Related concepts

Fast copy operations

Related tasks

Performing a fast copy operation

16

CIFS

This chapter includes:

Topics:

CIFS overview

Common Internet File System (CIFS) clients can have access to the system directories on the protection system.

- The `/data/col1/backup` directory is the destination directory for compressed backup server data.
- The `/ddvar/core` directory contains system core and log files (remove old logs and core files to free space in this area).

  NOTE: You can also delete core files from the `/ddvar` or the `/ddvar/ext` directory if it exists.

Clients, such as backup servers that perform backup and restore operations need access to the `/data/col1/backup` directory, at a minimum. Clients that have administrative access need to be able to access the `/ddvar/core` directory to retrieve core and log files.

As part of the initial protection system configuration, CIFS clients were configured to access these directories. This chapter describes how to modify these settings and how to manage data access using the DD System Manager and the `cifs` command.

NOTE:

- The DD System Manager **Protocols** > **CIFS** page allows you to perform major CIFS operations such as enabling and disabling CIFS, setting authentication, managing shares, and viewing configuration and share information.
- The `cifs` command contains all the options to manage CIFS backup and restores between Windows clients and protection systems, and to display CIFS statistics and status. For complete information about the `cifs` command, see the *DDOS Command Reference Guide*.
- For information about setting up clients to use the protection system as a server, see the related tuning guide, such as the *CIFS Tuning Guide*, which is available from https://www.dell.com/support. Search for the complete name of the document using the Search field.

Performing CIFS setup

This section contains instructions about enabling CIFS services, naming the CIFS server, and so on.

HA systems and CIFS

HA systems are compatible with CIFS; however, if a CIFS job is in progress during a failover, the job will need to be restarted.

"/ddvar is an ext3 file system, and cannot be shared like a normal MTree-based share. The information in /ddvar will become stale when the active node fails over to the standby node because the filehandles are different on the two nodes. If /ddvar is mounted to access log files or upgrade the system, unmount and remount /ddvar if a failover has occurred since the last time /ddvar was mounted."

Enabling CIFS services

Enable the client to access the system using the CIFS protocol.

About this task

After configuring a client for access to protection systems, enable CIFS services, which allows the client to access the system using the CIFS protocol.

Steps

1. For the system selected in the DD System Manager Navigation tree, click **Protocols** > **CIFS**.
2. In the CIFS Status area, click **Enable**.

Naming the CIFS server

The hostname for the protection system that serves as the CIFS server is set during the system's initial configuration.

To change a CIFS server name, see the procedures in the section regarding setting authentication parameters.

A system's hostname should match the name assigned to its IP address, or addresses, in the DNS table. Otherwise authentication, as well as attempts to join a domain, can fail. If you need to change the system's hostname, use the `net set hostname` command, and also modify the system's entry in the DNS table.

When the system acts as a CIFS server, it takes the hostname of the system. For compatibility purposes, it also creates a NetBIOS name. The NetBIOS name is the first component of the hostname in all uppercase letters. For example, the hostname `jp9.oasis.local` is truncated to the NetBIOS name `JP9`. The CIFS server responds to both names.

You can have the CIFS server respond to different names at the NetBIOS levels by changing the NetBIOS hostname.

Related concepts

Setting authentication parameters

Changing the NetBIOS hostname

Change the NetBIOS hostname with the CLI.

Steps

1. Display the current NetBIOS name by entering:
   `# cifs show config`
2. Use the `cifs set nb-hostname` *`nb-hostname`* command.

Setting authentication parameters

Set the authentication parameters for working with CIFS.

Click the Configure link in to the left of the Authentication label in the Configuration tab. The system will navigate to the **Administration** > **Access** > **Authentication** tab where you can configure authentication for Active Directory, Kerberos, Workgroups, and NIS.

Related concepts

Naming the CIFS server

Setting CIFS options

View CIFS configuration, restrict anonymous connections.

Steps

1. Select **Protocols** > **CIFS** > **Configuration**.
2. In the Options area, click **Configure Options**.
3. To restrict anonymous connections, click the checkbox of the **Enable** option in the Restrict Anonymous Connections area.
4. In the Log Level area, click the drop-down list to select the level number.

   The level is an integer from 1 (one) to 5 (five). One is the default system level that sends the least-detailed level of CIFS-related log messages, five results in the most detail. Log messages are stored in the file `/ddvar/log/debug/cifs/cifs.log`.

   NOTE: A log level of 5 degrades system performance. Click the **Default** in the Log Level area after debugging an issue. This sets the level back to 1.

5. In the Server Signing area, select:
   - **Enabled** to enable server signing
   - **Disabled** to disable server signing
   - **Required** when server signing is required
6. Enable support for SMBv1 and SMBv2 as required.

   SMBv1 is disabled by default. This limitation applies to both fresh installations and upgrades from previous releases. Select it in the GUI or run the `cifs option set support-smb1 enabled`command to enable it if required.

   NOTE: If SMBv1 is required, run this command to enable it before initiating a DDOS upgrade.

Related references

Configuring SMB signing

Disabling CIFS services

Prevent clients from accessing the protection system.

Steps

1. Select **Protocols** > **CIFS**.
2. In the Status area, click **Disable**.
3. Click **OK**.

   Even after disabling CIFS access, CIFS authentication services continue to run on the system. This continuation is required to authenticate active directory domain users for management access.

Working with shares

To share data, create shares on the protection system.

Shares are administered on the protection system and the CIFS systems.

Creating shares

When creating shares, you have to assign client access to each directory separately and remove access from each directory separately. For example, a client can be removed from `/ddvar` and still have access to `/data/col1/backup`

About this task

A protection system supports a maximum number of 3000 CIFS shares,[1] and 600 simultaneous connections are allowed. However, the maximum number of connections that are supported is based on system memory. See the section regarding setting the maximum open files on a connection for more information.

NOTE: If Replication is to be implemented, a system can receive backups from both CIFS clients and NFS clients as long as separate directories are used for each. Do not mix CIFS and NFS data in the same directory.

Do not use the top level of an MTrree to host a CIFS share. Create a subdirectory within the MTree, and specify that subdirectory as the path for the CIFS share.

Steps

1. Select **Protocols** > **CIFS** tabs to go to the CIFS view.
2. Ensure that authentication has been configured, as described in the section regarding setting authentication parameters.
3. On the CIFS client, set shared directory permissions or security options.
4. On the CIFS view, click the Shares tab.
5. Click **Create**.
6. In the Create Shares dialog box, enter the following information:

Table 11. Shares dialog box information

Item	Description
Share Name	A descriptive name for the share.
Directory Path	The path to the target directory (for example, `/data/col1/backup/dir1`). **NOTE:** col1 uses the lower case letter L followed by the number 1.
Comment	A descriptive comment about the share.

NOTE: The share name can be a maximum of 80 characters and cannot contain the following characters: \ / : * ? " < > | + [ ] ; , = or extended ASCII characters.

7. Add a client by clicking Add (**+**) in the Clients area. The Client dialog box is displayed. Enter the name of the client in the Client text box and click **OK**.

   Consider the following when entering the client name.
   - No blanks or tabs (white space) characters are enabled.
   - It is not recommended to use both an asterisk (*) and individual client name or IP address for a given share. When an asterisk (*) is present, any other client entries for that share are not used.
   - It is not required to use both client name and client IP address for the same client on a given share. Use client names when the client names are defined in the DNS table.
   - To make share available to all clients, specify an asterisk (*) as the client. All users in the client list can access the share, unless one or more user names are specified, in which case only the listed names can access the share.

   Repeat this step for each client that you need to configure.
8. In the Max Connections area, select the text box and enter the maximum number of connections to the share that are enabled at one time. The default value of zero (also settable through the Unlimited button) enforces no limit on the number of connections.
9. Click **OK**.

   The newly created share is displayed at the end of the list of shares, which are located in the center of the Shares panel.

[1] May be affected by hardware limitations.

Related concepts

Setting authentication parameters

CLI equivalent

Steps

1. Run the `cifs status` command to verify that CIFS is enabled.
2. Run the `filesys status` command to verify that file system is enabled.
3. Run the `hostname` command to determine the system hostname.
4. Create the CIFS share.

   `cifs share create <share> path <path> {max-connections <max connections> | clients <clients> | users <users> | comment <comment>}`

   ```
   # cifs share create backup path /backup
   ```

5. Grant client access to the share.

   `cifs share modify <share> {max-connections <max connections> | clients <clients> | users <users> | comment <comment>}`

   ```
   # cifs share modify backup clients "srvr24.yourdomain.com,srvr24,10.24.160.116
   ```

6. From the Windows system, select **Start** > **Run**, and type the hostname and directory of the CIFS share.

   \\<DDhostname>.<DDdomain.com>\<sharename>

7. If there are problems connecting to the CIFS share, run the `cifs share show` command to verify the status of the share.

   The warning `WARNING: The share path does not exist!` is displayed if the share does not exist or was misspelled on creation.

   ```
   # cifs share show
   --------------- share backup ---------------

   enabled: yes
   path: /backup
   ```

8. If the CIFS share is still not accessible, verify that all client information is in the access list, and all network connections are functional.

Modifying a share

Change share information and connections.

Steps

1. Select **Protocols** > **CIFS** > **Shares** to navigate to the CIFS view, Shares tab.
2. Click the checkbox next the share that you wish to modify in the Share Name list.
3. Click **Modify**.
4. Modify share information:
   a. To change the comment, enter new text in the Comment text field.
   b. To modify a User or Group names, in the User/Group list, click the checkbox of the user or group and click **Edit** (pencil icon) or **Delete** (X). To add a user or group, click (**+**), and in the User/Group dialog box select the Type for User or Group, and enter the user or group name.
   c. To modify a client name, in the Client list click the checkbox of the client and click **Edit** (pencil icon) or **Delete** (X). To add a client, click the Add (**+**) and add the name in the Client dialog box.

      NOTE: To make the share available to all clients, specify an asterisk (*) as the client. All users in the client list can access the share, unless one or more user names are specified, in which case only the listed names can access the share.

   d. Click **OK**.

5. In the Max Connections area, in the text box, change the maximum number of connections to the share that are allowed at one time. Or select Unlimited to enforce no limit on the number of connections.
6. Click **OK**.

Creating a share from an existing share

Create a share from an existing share and modify the new share if necessary.

About this task

NOTE: User permissions from the existing share are carried over to the new share.

Steps

1. In the CIFS Shares tab, click the checkbox for the share you wish to use as the source.
2. Click **Create From**.
3. Modify the share information, as described in the section about modifying a share.

Related tasks

Modifying a share

Disabling a share

Disable one or more existing shares.

Steps

1. In the Shares tab, click the checkbox of the share you wish to disable in the Share Name list.
2. Click **Disable**.
3. Click **Close**.

Enabling a share

Enable one or more existing shares.

Steps

1. In the Shares tab, click the checkbox of the shares you wish to enable in the Share Name list.
2. Click **Enable**.
3. Click **Close**.

Deleting a share

Delete one or more existing shares.

Steps

1. In the Shares tab, click the checkbox of the shares you wish to delete in the Share Name list.
2. Click **Delete**.

   The Warning dialog box appears.

3. Click **OK**.

   The shares are removed.

Performing MMC administration

Use the Microsoft Management Console (MMC) for administration.

DDOS supports these MMC features:

- Share management, except for browsing when adding a share, or the changing of the offline settings default, which is a manual procedure.
- Session management.
- Open file management, except for deleting files.

Connecting to a protection system from a CIFS client

Use CIFS to connect to a protection system and create a read-only backup subfolder.

Steps

1. On the system CIFS page, verify that CIFS Status shows that CIFS is enabled and running.
2. In the Control Panel, open Administrative Tools and select **Computer Management**.
3. In the Computer Management dialog box, right-click **Computer Management (Local)** and select **Connect to another computer** from the menu.
4. In the Select Computer dialog box, select **Another computer** and enter the name or IP address for the protection system.
5. Create a `\backup` subfolder as read-only. For more information, see the section on creating a /data/col1/backup subfolder as read-only.

Creating a \data\col1\backup subfolder as read-only

Enter a path, share name, and select permissions.

Steps

1. Right-click **Shares** in the Shared Folders directory.
2. Select **New File Share** from the menu.

   The **Create a Shared Folder** wizard opens. The computer name should be the name or IP address of the protection system.

3. Enter the path for the Folder to share, for example, enter `C:\data\col1\backup\newshare`.
4. Enter the Share name, for example, enter **`newshare`**. Click **Next**.
5. For the Share Folder Permissions, selected Administrators have full access. Other users have read-only access. Click **Next**.
6. The Completing dialog box shows that you have successfully shared the folder with all Microsoft Windows clients in the network. Click **Finish**.

   The newly created shared folder is listed in the Computer Management dialog box.

Displaying CIFS information

Display information about shared folders, sessions, and open files.

Steps

1. In the Control Panel, open Administrative Tools and select **Computer Management**.
2. Select one of the Shared Folders (**Shares**, **Sessions**, or **Open Files**) in the System Tools directory.

   Information about shared folders, sessions, and open files is shown in the right panel.

Configuring SMB signing

On a DDOS version that supports it, you can configure the SMB signing feature using the CIFS option called server signing.

This feature is disabled by default because it degrades performance. When enabled, SMB signing can cause a 29 percent (reads) to 50 percent (writes) throughput performance drop, although individual system performance will vary. There are three possible values for SMB signing: disabled, auto and mandatory:

- When SMB signing is set to disabled, SMB signing is disabled, this is the default.
- When SMB signing is set to required, SMB signing is required, and both computers in the SMB connection must have SMB signing enabled.

SMB Signing CLI Commands

`cifs option set "server-signing" required`

Sets server signing to required.

`cifs option reset "server-signing"`

Resets server signing to the default (disabled).

As a best practice, whenever you change the SMB signing options, disable and then enable (restart) CIFS service using the following CLI commands:

`cifs disable`

`cifs enable`

The DD System Manager interface displays whether the SMB signing option is disabled or set to auto or mandatory. To view this setting in the interface, navigate to: **Protocols** > **CIFS** > **Configuration tab**. In the Options area, the value for the SMB signing option will be disabled, auto or mandatory reflecting the value set using the CLI commands.

Managing access control

Access shared from a Windows client, provide administrative access, and allow access from trusted domain users.

Accessing shares from a Windows client

Use the command line to map a share.

Steps

- From the Windows client use this DOS command: `net use` *`drive`*`:` *`backup-location`*

Example

For example, enter:

`# \\PP02\backup /USER:PP02\backup22`

This command maps the backup share from PowerProtect system PP02 to drive H on the Windows system and gives the user named backup22 access to the `\\PP_sys\backup` directory.

DDOS supports the SMB Change Notify functionality. This improves CIFS performance on the Windows client by allowing the CIFS server to automatically notify the Windows client about changes on the CIFS share, and eliminate the need for the client to poll the protection system to look for changes to the share.

Providing domain users administrative access

Use the command line to add CIFS and include the domain name in the ssh instruction.

Steps

- Enter: **`adminaccess authentication add cifs`**

  The SSH, Telnet, or FTP command that accesses the protection system must include, in double quotation marks, the domain name, a backslash, and the user name. For example:

  `C:> ssh "domain2\djones" @dd22`

Allowing administrative access to a protection system for domain users

Use the command line to map a DD system default group number, and then enable CIFS administrative access.

Steps

1. To map a protection system default group number to a Windows group name that differs from the default group name, use the `cifs option set "dd admin group2" ["windows grp-name"]` command.

   The Windows group name is a group (based on one of the user roles—admin, user, or back-up operator) that exists on a Windows domain controller, and you can have up to 50 groups (dd admin group1 to dd admin group50).

   NOTE: For a description of DDOS user roles and Windows groups, see the section about managing protection systems.

2. Enable CIFS administrative access by entering:
   `adminaccess authentication add cifs`
   - The default system group dd admin group1 is mapped to the Windows group Domain Admins.
   - You can map the default system group dd admin group2 to a Windows group named Data Domain that you create on a Windows domain controller.
   - Access is available through SSH, Telnet, FTP, HTTP, and HTTPS.
   - After setting up administrative access to the protection system from the Windows group **`Data Domain`**, you must enable CIFS administrative access using the `adminaccess` command.

Restricting administrative access from Windows

Use the command line to prohibit access to users without a DD account.

Steps

- Enter: **`adminaccess authentication del cifs`**

  This command prohibits Windows users access to the protection system if they do not have an account on the system.

File access

This sections contains information about ACLs, setting DACL and SACL permissions using Windows Explorer, and so on.

NT access control lists

Access control lists (ACLs) are enabled by default on the protection system.

CAUTION: Dell recommends that you do not disable NTFS ACLs once they have been enabled. Contact Dell Support prior to disabling NTFS ACLs.

Default ACL Permissions

The default permissions, which are assigned to new objects created through the CIFS protocol when ACLs are enabled, depend on the status of the parent directory. There are three different possibilities:

- The parent directory has no ACL because it was created through NFS protocol.
- The parent directory has an inheritable ACL, either because it was created through the CIFS protocol or because ACL had been explicitly set. The inherited ACL is set on new objects.
- The parent directory has an ACL, but it is not inheritable. The permissions are as follows:

Table 12. Permissions

Type	Name	Permission	Apply To
Allow	SYSTEM	Full control	This folder only
Allow	CREATOR OWNER	Full control	This folder only

NOTE: CREATOR OWNER is replaced by the user creating the file/folder for normal users and by Administrators for administrative users.

Permissions for a New Object when the Parent Directory Has No ACL

The permissions are as follows:

- BUILTIN\Administrators:(OI)(CI)F
- NT AUTHORITY\SYSTEM:(OI)(CI)F
- CREATOR OWNER:(OI)(CI)(IO)F
- BUILTIN\Users:(OI)(CI)R
- BUILTIN\Users:(CI)(special access:)FILE_APPEND_DATA
- BUILTIN\Users:(CI)(IO)(special access:)FILE_WRITE_DATA
- Everyone:(OI)(CI)R

These permissions are described in more detail as follows:

Table 13. Permissions Detail

Type	Name	Permission	Apply To
Allow	Administrators	Full control	This folder, subfolders, and files
Allow	SYSTEM	Full control	This folder, subfolders, and files
Allow	CREATOR OWNER	Full control	Subfolders and files only
Allow	Users	Read & execute	This folder, subfolders, and files
Allow	Users	Create subfolders	This folder and subfolders only
Allow	Users	Create files	Subfolders only
Allow	Everyone	Read & execute	This folder, subfolders, and files

Setting ACL Permissions and Security

Windows-based backup and restore tools such as NetBackup can be used to back up DACL- and SACL-protected files to, and restore them from, the protection system.

Granular and Complex Permissions (DACL)

You can set granular and complex permissions (DACL) on any file or folder object within the file system, either through using Windows commands such as `cacls`, `xcacls`, `xcopy` and `scopy`, or through the CIFS protocol using the Windows Explorer GUI.

Audit ACL (SACL)

You can set audit ACL (SACL) on any object in the file system, either through commands or through the CIFS protocol using the Windows Explorer GUI.

Setting DACL permissions using the Windows Explorer

Use Explorer properties settings to select DACL permissions.

Steps

1. Right-click the file or folder and select **Properties**.
2. In the Properties dialog box, click the Security tab.
3. Select the group or user name, such as **Administrators**, from the list. The permissions appear, in this case for `Administrators, Full Control`.
4. Click the **Advanced** button, which enables you to set special permissions.
5. In the Advanced Security Settings for ACL dialog box, click the Permissions tab.
6. Select the permission entry in the list.
7. To view more information about a permission entry, select the entry and click **Edit**.
8. Select the Inherit from parent option to have the permissions of parent entries inherited by their child objects, and click **OK**.

Setting SACL permissions using the Windows Explorer

Use Explorer properties settings to select SACL permissions.

Steps

1. Right-click the file or folder and select **Properties** from the menu.
2. In the Properties dialog box, click the Security tab.
3. Select the group or user name, such as **Administrators**, from the list, which displays its permissions, in this case, `Full Control`.
4. Click the **Advanced** button, which enables you to set special permissions.
5. In the Advanced Security Settings for ACL dialog box, click the Auditing tab.
6. Select the auditing entry in the list.
7. To view more information about special auditing entries, select the entry and click **Edit**.
8. Select the Inherit from parent option to have the permissions of parent entries inherited by their child objects, and click **OK**.

Viewing or changing the current owner security ID (owner SID)

Use the Advanced Security Settings for ACL dialog box.

Steps

1. In the Advanced Security Settings for ACL dialog box, click the Owner tab.
2. To change the owner, select a name from the Change owner list, and click **OK**.

Controlling ID account mapping

The CIFS option idmap-type controls ID account mapping behavior.

This option has two values: rid (the default) and none. When the option is set to rid, the ID-to-id mapping is performed internally. When the option is set to none, all CIFS users are mapped to a local UNIX user named "cifsuser" belonging to the local UNIX group users.

Consider the following information while managing this option.

- CIFS must be disabled to set this option. If CIFS is running, disable CIFS services.

- The idmap-type can be set to none only when ACL support is enabled.
- Whenever the idmap type is changed, a file system metadata conversion might be required for correct file access. Without any conversion, the user might not be able to access the data. To convert the metadata, consult your contracted support provider.

Monitoring CIFS operation

Monitoring CIFS Operation topics.

Displaying CIFS status

View and enable/disable CIFS status.

Steps

1. In the DD System Manager, select **Protocols** > **CIFS**.
   - Status is either enabled and running, or disabled but CIFS authentication is running.

     To enable CIFS, see the section regarding enabling CIFS services. To disable CIFS, see the section regarding disabling CIFS services.
   - **Connections** lists the tally of open connections and open files.

   Table 14. Connections Details information

Item	Description
Open Connections	Open CIFS connections
Connection Limit	Maximum allowed connections
Open Files	Current open files
Max Open Files	Maximum number of open files

2. Click **Connection Details** to see more connection information.

   Table 15. Connections Details information

Item	Description
Sessions	Active CIFS sessions
Computer	IP address or computer name connected with DDR for the session
User	User operating the computer connected with the DDR
Open Files	Number of open files for each session
Connection Time	Connection length in minutes
User	Domain name of computer
Mode	File permissions
Locks	Number of locks on the file
Files	File location

Related concepts

Display shares information

Related tasks

Enabling CIFS services
Disabling CIFS services

Display CIFS configuration

This section displays CIFS Configuration.

Authentication configuration

The information in the Authentication panel changes, depending on the type of authentication that is configured.

Click the Configure link in to the left of the Authentication label in the Configuration tab. The system will navigate to the **Administration** > **Access** > **Authentication** page where you can configure authentication for Active Directory, Kerberos, Workgroups, and NIS.

Active directory configuration

Table 16. Active directory configuration information

Item	Description
Mode	The Active Directory mode displays.
Realm	The configured realm displays.
DDNS	The status of the DDNS Server displays: either enabled or disabled.
Domain Controllers	The name of the configured domain controllers display or a * if all controllers are permitted.
Organizational Unit	The name of the configured organizational units displays.
CIFS Server Name	The name of the configured CIFS server displays.
WINS Server Name	The name of the configured WINS server displays.
Short Domain Name	The short domain name displays.

Workgroup configuration

Table 17. Workgroup configuration authentication information

Item	Description
Mode	The Workgroup mode displays.
Workgroup Name	The configured workgroup name displays.
DDNS	The status of the DDNS Server displays: either enabled or disabled.
CIFS Server Name	The name of the configured CIFS server displays.
WINS Server Name	The name of the configured WINS server displays.

Display shares information

This section displays shares information.

Viewing configured shares

View the list of configured shares.

Table 18. Configured shares information

Item	Description
Share Name	The name of the share (for example, share1).
Share Status	The status of the share: either enabled or disabled.

Table 18. Configured shares information (continued)

Item	Description
Directory Path	The directory path to the share (for example, `/data/col1/backup/dir1`). **NOTE:** col1 uses the lower case letter L followed by the number 1.
Directory Path Status	The status of the directory path.

- To list information about a specific share, enter the share name in the Filter by Share Name text box and click **Update**.
- Click **Update** to return to the default list.
- To page through the list of shares, click the **<** and **>** arrows at the bottom right of the view to page forward or backward. To skip to the beginning of the list, click **|<** and to skip to the end, click **>|**.
- Click the **Items per Page** drop-down arrow to change the number of share entries listed on a page. Choices are 15, 30, or 45 entries.

Viewing detailed share information

Display detailed information about a share by clicking a share name in the share list.

Table 19. Share information

Item	Description
Share Name	The name of the share (for example, share1).
Directory Path	The directory path to the share (for example, `/data/col1/backup/dir1`). **NOTE:** col1 uses the lower case letter L followed by the number 1.
Directory Path Status	Indicates whether the configured directory path exists on the DDR. Possible values are Path Exists or Path Does Not Exist, the later indicating an incorrect or incomplete CIFS configuration.
Max Connections	The maximum number of connections allowed to the share at one time. The default value is Unlimited.
Comment	The comment that was configured when the share was created.
Share Status	The status of the share: either enabled or disabled.

- The Clients area lists the clients that are configured to access the share, along with a client tally beneath the list.
- The User/Groups area lists the names and type of users or groups that are configured to access the share, along with a user or group tally beneath the list.
- The Options area lists the name and value of configured options.

Displaying CIFS statistics

Use the command line to display CIFS statistics.

Steps

- Enter: **`cifs show detailed-stats`**

  The output shows number of various SMB requests received and the time taken to process them.

Performing CIFS troubleshooting

This section provides basic troubleshooting procedures.

NOTE: The `cifs troubleshooting` commands provide detailed information about CIFS users and groups.

Displaying clients current activity

Use the command line to display CIFS sessions and open files information.

Steps

- Enter: **`cifs show active`**

Results

Table 20. Sessions

Computer	User	Open files	Connect time (sec)	Idle time (sec)
::ffff:10.25.132.84	ddve-25179109\sysadmin	1	92	0

Table 21. Open files

User	Mode	Locks	File
ddve-25179109\sysadmin	1	0	C:\data\col1\backup

Setting the maximum open files on a connection

Use the command line to set the maximum number of files that can be open concurrently.

Steps

- Enter: **`cifs option set max-global-open-files`** *value*.

  The *value* for the maximum global open files can be between 1 and the open files maximum limit. The maximum limit is based on the DDR system memory. For systems with greater than 12 GB, the maximum open files limit is 30,000. For systems with less than or equal to 12 GB, the maximum open files limit is 10,000.

  Table 22. Connection and maximum open file limits

Memory	Connection Limit	Open File Maximum Limit
8 GB	300	10,000
16 GB and higher	600	30,000

  NOTE: The system has a maximum limit of 600 CIFS connections and 250,000 open files. However, if the system runs out of open files, the number of files can be increased.

  NOTE: File access latencies are affected by the number of files in a directory. To the extent possible, we recommend directory sizes of less than 250,000. Larger directory sizes might experience slower responses to metadata operations such as listing the files in the directory and opening or creating a file.

System clock

When using active directory mode for CIFS access, the system clock time can differ by no more than five minutes from that of the domain controller.

When configured for Active Directory authentication, the system regularly syncs time with the Windows domain controller. Therefore, it is important for the domain controller to obtain the time from a reliable time source. Refer to the Microsoft documentation for your Windows operating system version to configure the domain controller with a time source.

WARNING: When the system is configured for Active Directory authentication, it uses an alternate mechanism to sync time with the domain controller. To avoid time sync conflicts, do not enable NTP when the system is configured for Active Directory authentication.

Synchronize from an NTP server

Configure the time server synchronization, as described in the section regarding working with time and date settings.

17

NFS

This chapter includes:

Topics:

NFS overview

Network File System (NFS) clients can have access to the system directories or MTrees on the protection system.

- The `/backup` directory is the default destination for non-MTree compressed backup server data.
- The `/data/col1/backup` path is the root destination when using MTrees for compressed backup server data.
- The `/ddvar/core` directory contains system core and log files (remove old logs and core files to free space in this area).

  NOTE: On protection systems, the `/ddvar/core` is on a separate partition. If you mount `/ddvar` only, you will not be able to navigate to `/ddvar/core` from the `/ddvar` mountpoint.

Clients, such as backup servers that perform backup and restore operations need access to the `/backup` or `/data/col1/backup` areas, at a minimum. Clients that have administrative access need to be able to access the `/ddvar/core` directory to retrieve core and log files.

As part of the initial system configuration, NFS clients were configured to access these areas. This chapter describes how to modify these settings and how to manage data access.

NOTE:

- The `nfs` command manages backups and restores between NFS clients and protection systems, and it displays NFS statistics and status. For complete information about the `nfs` command, see the *DDOS Command Reference Guide*.
- For information about setting up third-party clients to use the protection system as a server, see the related tuning guide, such as the *Solaris System Tuning*, which is available from the Dell Support web site.

HA systems and NFS

HA systems are compatible with NFS. If a NFS job is in progress during a failover, the job will **not** need to be restarted.

NOTE: /ddvar is an ext3 file system, and cannot be shared like a normal MTree-based share. The information in /ddvar will become stale when the active node fails over to the standby node because the filehandles are different on the two nodes. If /ddvar is mounted to access log files or upgrade the system, unmount and remount /ddvar if a failover has occurred since the last time /ddvar was mounted.

To create valid NFS exports that will failover with HA, the export needs to be created from the Active HA node, and generally shared over the failover network interfaces.

Managing NFS client access to the protection system

The topics in this section describe how to manage NFS client access to a protection System.

The KB article *NFS Best Practices for Data Domain and client OS*, available from the Online Support website, provides additional information about best practices for NFS.

Enabling NFS services

Enable NFS services to allow the client to access the system using the NFS protocol.

Steps

1. Select **Protocols** > **NFS**.

   The NFS view opens displaying the Exports tab.

2. Click **Enable**.

Disabling NFS services

Disable NFS services to prevent the client access to the system using the NFS protocol.

Steps

1. Select the **Protocols** > **NFS** tabs.

   The NFS view opens displaying the Exports tab.

2. Click **Disable**.

Creating an export

You can use DD SM's Create button on the NFS view or use the Configuration Wizard to specify the NFS clients that can access the `/backup`, `/data/col1/backup`,`/ddvar`, `/ddvar/core` areas, or the`/ddvar/ext` area if it exists.

About this task

A protection system supports a maximum of 2048 exports[2], with the number of connections scaling in accordance with system memory.

NOTE: You have to assign client access to each export separately and remove access from each export separately. For example, a client can be removed from `/ddvar` and still have access to `/data/col1/backup`.

CAUTION: If Replication is to be implemented, a single destination system can receive backups from both CIFS clients and NFS clients as long as separate directories or MTrees are used for each. Do not mix CIFS and NFS data in the same area.

Do not use the top level of an MTrree to host an NFS export. Create a subdirectory within the MTree, and specify that subdirectory as the path for the NFS export.

Steps

1. Select **ProtocolsNFS**.

   The NFS view opens displaying the Exports tab.

2. Click **Create**.
3. Enter the pathname in the Directory Path text box (for example, `/data/col1/backup/dir1`).

   NOTE: `col1` uses the lower-case letter L followed by the number 1.

4. In the Clients area, select an existing client or click the **+** icon to create a client.

   The Client dialog box is displayed.

   a. Enter a server name in the text box.

      Enter fully qualified domain names, hostnames, or IP addresses. A single asterisk (*) as a wild card indicates that all backup servers are to be used as clients.

2 May be affected by hardware limitations.

NOTE: Clients given access to the `/data/col1/backup` directory have access to the entire directory. A client given access to a subdirectory of `/data/col1/backup` has access only to that subdirectory.

- A client can be a fully-qualified domain hostname, an IPv4 or IPv6 IP address, an IPv4 address with either a netmask or prefix length, an IPv6 address with prefix length, an NIS netgroup name with the prefix `@`, or an asterisk (`*`) wildcard with a domain name, such as `*.yourcompany.com`.
- A client added to a subdirectory under `/data/col1/backup` has access only to that subdirectory.
- Enter an asterisk (`*`) as the client list to give access to all clients on the network.

b. Select the checkboxes of the NFS options for the client.

General:

- Read-only permission (ro).
- Allow connections from ports below 1024 (secure) (default).

Anonymous UID/GID:

- Map requests from UID (user identifier) or GID (group identifier) 0 to the anonymous UID/GID (root _squash).
- Map all user requests to the anonymous UID/GID (all _squash).
- Use Default Anonymous UID/GID.

Allowed Kerberos Authentication Modes:

- Unauthenticated connections (sec=sys). Select to not use authentication.
- Authenticated Connections (sec=krb5).

NOTE: Integrity and Privacy are supported, although they might slow performance considerably.

c. Click **OK**.

5. Click **OK** to create the export.

Modifying an export

Change the directory path, domain name, and other options using the GUI.

Steps

1. Select**Protocols** > **NFS**.

   The NFS view opens displaying the Exports tab.

2. Click the checkbox of an export in the NFS Exports table.
3. Click **Modify**.
4. Modify the pathname in the Directory Path text box.
5. In the Clients area, select another client and click the pencil icon (modify), or click the **+** icon to create a client.

a. Enter a server name in the Client text box.

Enter fully qualified domain names, hostnames, or IP addresses. A single asterisk (*) as a wild card indicates that all backup servers are to be used as clients.

NOTE: Clients given access to the `/data/col1/backup` directory have access to the entire directory. A client given access to a subdirectory of `/data/col1/backup` has access only to that subdirectory.

- A client can be a fully-qualified domain hostname, an IPv4 or IPv6 IP address, an IPv4 address with either a netmask or prefix length, an IPv6 address with prefix length, an NIS netgroup name with the prefix `@`, or an asterisk (*) wildcard with a domain name, such as `*.yourcompany.com`.

  A client added to a subdirectory under `/data/col1/backup` has access only to that subdirectory.

- Enter an asterisk (`*`) as the client list to give access to all clients on the network.

b. Select the checkboxes of the NFS options for the client.

General:

- Read-only permission (ro).
- Allow connections from ports below 1024 (secure) (default).

Anonymous UID/GID:

- Map requests from UID (user identifier) or GID (group identifier) 0 to the anonymous UID/GID (root _squash).

- Map all user requests to the anonymous UID/GID (all _squash).
- Use Default Anonymous UID/GID.

Allowed Kerberos Authentication Modes:

- Unauthenticated connections (sec=sys). Select to not use authentication.
- Authenticated Connections (sec=krb5).

NOTE: Integrity and Privacy are not supported.

c. Click **OK**.

6. Click **OK** to modify the export.

Related tasks

Creating an export from an existing export

Creating an export from an existing export

Create an export from an existing export and then modify it as needed.

Steps

1. In the NFS Exports tab, click the checkbox of the export you wish to use as the source.
2. Click **Create From**.
3. Modify the export information, as described in section about modifying an export.

Related tasks

Modifying an export

Deleting an export

Delete an export from the NFS Exports tab.

Steps

1. In the NFS Exports tab, click the checkbox of the export you wish to delete.
2. Click **Delete**.
3. Click **OK** and **Close** to delete the export.

Displaying NFS information

The topics in this section describe how to use the DD System Manager to monitor NFS client status and NFS configuration.

Viewing NFS status

Display whether NFS is active and Kerberos is enabled.

Steps

- Click **Protocols** > **NFS**.

  The top panel shows the operational status of NFS; for example, whether NFS is currently active and running, and whether Kerberos mode is enabled.

  NOTE: Click Configure to view the **Administration** > **Access** > **Authentication** tab where you can configure Kerberos authentication.

Viewing NFS exports

See the list of clients allowed to access the protection system.

Steps

1. Click **Protocols** > **NFS**.

   The Exports view shows a table of NFS exports that are configured for system and the mount path, status, and NFS options for each export.

2. Click an export in the table to populate the Detailed Information area, below the Exports table.

   In addition to the export's directory path, configured options, and status, the system displays a list of clients.

   Use the Filter By text box to sort by mount path.

   Click **Update** for the system to refresh the table and use the filters supplied.

   Click **Reset** for the system to clear the Path and Client filters.

Viewing active NFS clients

Display all clients that have been connected in the past 15 minutes and their mount path.

Steps

- Select the **Protocols** > **NFS** > **Active Clients** tab.

  The Active Clients view displays, showing all clients that have been connected in the past 15 minutes and their mount path.

  Use the Filter By text boxes to sort by mount path and client name.

  Click **Update** for the system to refresh the table and use the filters supplied.

  Click **Reset** for the system to clear the Path and Client filters.

Integrating a DDR into a Kerberos domain

Set the domain name, the host name, and the DNS server for the DDR.

About this task

Enable the DDR to use the authentication server as a Key Distribution Center (for UNIX) and as a Distribution Center (for Windows Active Directory).

CAUTION: The examples provided in this description are specific to the operating system (OS) used to develop this exercise. You must use commands specific to your OS.

NOTE: For UNIX Kerberos mode, a keytab file must be transferred from the Key Distribution Center (KDC) server, where it is generated, to the DDR. If you are using more than one DDR, each DDR requires a separate keytab file. The keytab file contains a shared secret between the KDC server and the DDR.

NOTE: When using a UNIX KDC, the DNS server does not have to be the KDC server, it can be a separate server.

Steps

1. Set the host name and the domain name for the DDR, using DDR commands.

   ```
   net set hostname <host>
   ```

   ```
   net set {domainname <local-domain-name>}
   ```

   NOTE: The host name is the name of the DDR.

2. Configure NFS principal (node) for the DDR on the Key Distribution Center (KDC).

Example:

addprinc nfs/hostname@realm

NOTE: Hostname is the name for the DDR.

3. Verify that there are nfs entries added as principals on the KDC.
   Example:

   listprincs

   ```
   nfs/hostname@realm
   ```

4. Add the DDR principal into a keytab file.
   Example:

   ktadd <keytab_file> nfs/hostname@realm

5. Verify that there is an nfs keytab file configured on the KDC.
   Example:

   klist -k <keytab_file>

   NOTE: The <keytab_file> is the keytab file used to configure keys in a previous step.

6. Copy the keytab file from the location where the keys for NFS DDR are generated to the DDR in the /ddvar/ directory.

   Table 23. Keytab destination

Copy file from:	Copy file to:
<keytab_file> (The keytab file configured in a previous step.)	/ddvar/

7. Set the realm on the DDR, using the following DDR command:
   authentication kerberos set realm <home realm> kdc-type <unix, windows.> kdcs <IP address of server>
8. When the kdc-type is UNIX, import the keytab file from /ddvar/ to /ddr/etc/, where the Kerberos configuration file expects it. Use the following DDR command to copy the file:
   authentication kerberos keytab import

   NOTE: This step is required only when the kdc-type is UNIX.

   Kerberos setup is now complete.
9. To add a NFS mount point to use Kerberos, use the nfs add command.
   See the *DDOS Command Reference Guide* for more information.
10. Add host, NFS and relevant user principals for each NFS client on the Key Distribution Center (KDC).
    Example: **listprincs**

    ```
    host/hostname@realm
    nfs/hostname@realm
    root/hostname@realm
    ```

11. For each NFS client, import all its principals into a keytab file on the client.
    Example:

    ktadd -k <keytab_file> host/hostname@realm

    ktadd -k <keytab_file> nfs/hostname@realm

Add and delete KDC servers after initial configuration

After you have integrated a DDR into a Kerberos domain, and thereby enabled the DDR to use the authentication server as a Key Distribution Center (for UNIX) and as a Distribution Center (for Windows Active Directory), you can use the following procedure to add or delete KDC servers.

Steps

1. Join the DDR to a Windows Active Directory (AD) server or a UNIX Key Distribution Center (KDC).

   authentication kerberos set realm <home-realm> kdc-type {windows [kdcs <kdc-list>] | unix kdcs <kdc-list>}

   Example: **authentication kerberos set realm krb5.test kdc-type unix kdcs nfskrb-kdc.krb5.test**

   This command joins the system to the krb5.test realm and enables Kerberos authentication for NFS clients.

   NOTE: A keytab generated on this KDC must exist on the DDR to authenticate using Kerberos.

2. Verify the Kerberos authentication configuration.

   authentication kerberos show config

   ```
   Home Realm:          krb5.test
   KDC List:            nfskrb-kdc.krb5.test
   KDC Type:            unix
   ```

3. Add a second KDC server.

   authentication kerberos set realm <home-realm> kdc-type {windows [kdcs <kdc-list>] | unix kdcs <kdc-list>}

   Example: **authentication kerberos set realm krb5.test kdc-type unix kdcs ostqa-sparc2.krb5.test nfskrb-kdc.krb5.test**

   NOTE: A keytab generated on this KDC must exist on the DDR to authenticate using Kerberos.

4. Verify that two KDC servers are added.

   authentication kerberos show config

   ```
   Home Realm:          krb5.test
   KDC List:            ostqa-sparc2.krb5.test, nfskrb-kdc.krb5.test
   KDC Type:            unix
   ```

5. Display the value for the Kerberos configuration key.

   reg show config.keberos

   ```
   config.kerberos.home_realm = krb5.test
   config.kerberos.home_realm.kdc1 = ostqa-sparc2.krb5.test
   config.kerberos.home_realm.kdc2 = nfskrb-kdc.krb5.test
   config.kerberos.kdc_count = 2
   config.kerberos.kdc_type = unix
   ```

6. Delete a KDC server.

   Delete a KDC server by using the **authentication kerberos set realm <home-realm> kdc-type {windows [kdcs <kdc-list>] | unix kdcs <kdc-list>}** command without listing the KDC server that you want to delete. For example, if the existing KDC servers are kdc1, kdc2, and kdc3, and you want to remove kdc2 from the realm, you could use the following example:

   authentication kerberos set realm <realm-name> kdc-type <kdc_type> kdcs kdc1,kdc3

18

NFSv4

This chapter includes:

Topics:

Introduction to NFSv4

Because NFS clients are increasingly using NFSv4.x as the default NFS protocol level, protection systems can now employ NFSv4 instead of requiring the client to work in a backwards-compatibility mode.

Clients can work in mixed environments in which NFSv4 and NFSv3 must be able to access the same NFS exports.

The DDOS NFS server can be configured to support NFSv4 and NFSv3, depending on site requirements. You can make each NFS export available to only NFSv4 clients, only NFSv3 clients, or both.

Several factors might affect whether you choose NFSv4 or NFSv3:

- NFS client support

  Some NFS clients may support only NFSv3 or NFSv4, or may operate better with one version.

- Operational requirements

  An enterprise might be strictly standardized to use either NFSv4 or NFSv3.

- Security

  If you require greater security, NFSv4 provides a greater security level than NFSv3, including ACL and extended owner and group configuration.

- Feature requirements

  If you need byte-range locking or UTF-8 files, you should choose NFSv4.

- NFSv3 submounts

  If your existing configuration uses NFSv3 submounts, NFSv3 might be the appropriate choice.

NFSv4 compared to NFSv3

NFSv4 provides enhanced functionality and features compared to NFSv3.

The following table compares NFSv3 features to those for NFSv4.

Table 24. NFSv4 compared to NFSv3

Feature	NFSv3	NFSv4
Standards-based Network Filesystem	Yes	Yes
Kerberos support	Yes	Yes
Kerberos with LDAP	Yes	Yes
Quota reporting	Yes	Yes
Multiple exports with client-based access lists	Yes	Yes
ID mapping	Yes	Yes
UTF-8 character support	No	Yes
File/directory-based Access Control Lists (ACL)	No	Yes
Extended owner/group (OWNER@)	No	Yes
File share locking	No	Yes
Byte range locking	No	Yes
DD-CIFS integration (locking, ACL, AD)	No	Yes
Stateful file opens and recovery	No	Yes
Global namespace and pseudoFS	No	Yes
Multi-system namespace using referrals	No	Yes

NFSv4 ports

You can enable or disable NFSv4 and NFSv3 independently. In addition, you can move NFS versions to different ports; both versions do not need to occupy the same port.

With NFSv4, you do not need to restart the file system if you change ports. Only an NFS restart is required in such instances.

Like NFSv3, NFSv4 runs on Port 2049 as the default if it is enabled.

NFSv4 does not use portmapper (Port 111) or mountd (Port 2052).

ID Mapping Overview

NFSv4 identifies owners and groups by a common external format, such as `joe@example.com`. These common formats are known as identifiers, or IDs.

Identifiers are stored within an NFS server and use internal representations such as ID 12345 or ID S-123-33-667-2. The conversion between internal and external identifiers is known as ID mapping.

Identifiers are associated with the following:

- Owners of files and directories
- Owner groups of files and directories
- Entries in Access Control Lists (ACLs)

Protection systems use a common internal format for NFS and CIFS/SMB protocols, which allows files and directories to be shared between NFS and CIFS/SMB. Each protocol converts the internal format to its own external format with its own ID mapping.

External formats

The external format for NFSv4 identifiers follows NFSv4 standards (for example, RFC-7530 for NFSv4.0). In addition, supplemental formats are supported for interoperability.

Standard identifier formats

Standard external identifiers for NFSv4 have the format *identifier@domain*. This identifier is used for NFSv4 owners, owner-groups, and access control entries (ACEs). The domain must match the configured NFSv4 domain that was set using the `nfs option` command.

The following CLI example sets the NFSv4 domain to `mycorp.com` for the NFS server:

```
nfs option set nfs4-domain mycorp.com
```

See client-specific documentation you have for setting the client NFS domain. Depending on the operating system, you might need to update a configuration file (for example, `/etc/idmapd.conf`) or use a client administrative tool.

NOTE: If you do not set the default value, it will follow the DNS name for the protection system.

NOTE: The file system must be restarted after changing the DNS domain for the nfs4-domain to automatically update.

ACE extended identifiers

For ACL ACE entries, protection system NFS servers also support the following standard NFSv4 ACE extended identifiers defined by the NFSv4 RFC:

- OWNER@, The current owner of the file or directory
- GROUP@, the current owner group of the file or directory.
- The special identifiers INTERACTIVE@, NETWORK@, DIALUP@, BATCH@, ANONYMOUS@, AUTHENTICATED@, SERVICE@.

Alternative formats

To allow interoperability, NFSv4 servers on protection systems support some alternative identifier formats for input and output.

- Numeric identifiers; for example, “12345”.
- Windows compatible Security identifiers (SIDs) expressed as “S-NNN-NNN-…”

See the sections on input mapping and output mapping for more information about restrictions to these formats.

Internal Identifier Formats

The DD file system stores identifiers with each object (file or directory) in the filesystem. All objects have a numeric user ID (UID) and group ID (GID). These, along with a set of mode bits, allow for traditional UNIX/Linux identification and access controls.

Objects created by the CIFS/SMB protocol, or by the NFSv4 protocol when NFSv4 ACLs are enabled, also have an extended security descriptor (SD). Each SD contains the following:

- An owner security identifier (SID)
- An owner group SID
- A discretionary ACL (DACL)
- (Optional) A system ACL (SACL)

Each SID contains a relative ID (RID) and a distinct domain in a similar manner to Windows SIDs. See the section on NFSv4 and CIFS interoperability for more information on SIDs and the mapping of SIDs.

When ID mapping occurs

The protection system NFSv4 server performs mapping in the following circumstances:

- Input mapping

  The NFS server receives an identifier from an NFSv4 client. See Input mapping.

- Output mapping:

  An identifier is sent from the NFS server to the NFSv4 client. See Output mapping.

- Credential mapping

  The RPC client credentials are mapped to an internal identity for access control and other operations. See Credential mapping.

Input mapping

Input mapping occurs when an NFSv4 client sends an identifier to the protection system NFSv4 server—setting up the owner or owner-group of a file, for example. Input mapping is distinct from credential mapping.

Standard format identifiers such as `joe@mycorp.com` are converted into an internal UID/GID based on the configured conversion rules. If NFSv4 ACLs are enabled, a SID will also be generated, based on the configured conversion rules.

Numeric identifiers (for example, "12345") are directly converted into corresponding UID/GIDs if the client is not using Kerberos authentication. If Kerberos is being used, an error will be generated as recommended by the NFSv4 standard. If NFSv4 ACLs are enabled, a SID will be generated based on the conversion rules.

Windows SIDs (for example, "S-NNN-NNN-…") are validated and directly converted into the corresponding SIDs. A UID/GID will be generated based on the conversion rules.

Output mapping

Output mapping occurs when the NFSv4 server sends an identifier to the NFSv4 client; for example, if the server returns the owner or owner-group of a file.

1. If configured, the output might be the numeric ID.

   This can be useful for NFSv4 clients that are not configured for ID mapping (for example, some Linux clients).

2. Mapping is attempted using the configured mapping services, (for example, NIS or Active Directory).
3. The output is a numeric ID or SID string if mapping fails and the configuration is allowed.
4. Otherwise, nobody is returned.

The **nfs option *nfs4-idmap-out-numeric*** configures the mapping on output:

- If **nfs option *nfs4-idmap-out-numeric*** is set to **map-first**, mapping will be attempted. On error, a numeric string is output if allowed. This is the default.
- If **nfs option *nfs4-idmap-out-numeric*** is set to **always**, output will always be a numeric string if allowed.
- If **nfs option *nfs4-idmap-out-numeric*** is set to **never**, mapping will be attempted. On error, **nobody@nfs4-domain** is the output.

  If the RPC connection uses GSS/Kerberos, a numeric string is never allowed and **nobody@nfs4-domain** is the output.

The following example configures the protection system NFS server to always attempt to output a numeric string on output. For Kerberos the name nobody is returned:

```
nfs option set nfs4-idmap-out-numeric always
```

Credential mapping

The NFSv4 server provides credentials for the NFSv4 client.

These credentials perform the following functions:

- Determine the access policy for the operation; for example, the ability to read a file.
- Determine the default owner and owner-group for new files and directories.

Credentials sent from the client may be `john_doe@mycorp.com`, or system credentials such as `UID=1000, GID=2000`. System credentials specify a UID/GID along with auxiliary group IDs.

If NFSv4 ACLs are disabled, then the UID/GID and auxiliary group IDs are used for the credentials.

If NFSv4 ACLs are enabled, then the configured mapping services are used to build an extended security descriptor for the credentials:

- SIDs for the owner, owner-group, and auxiliary group mapped and added to the Security Descriptor (SD).
- Credential privileges, if any, are added to the SD.

NFSv4 and CIFS/SMB Interoperability

The security descriptors used by NFSv4 and CIFS are similar from an ID mapping perspective, although there are differences.

You should be aware of the following to ensure for optimal interoperability:

- Active Directory should be configured for both CIFS and NFSv4, and the NFS ID mapper should be configured to use Active Directory for ID mapping.
- If you are using CIFS ACLs extensively, you can usually improve compatibility by also enabling NFSv4 ACLs.
  - Enabling NFSv4 ACLs allows NFSv4 credentials to be mapped to the appropriate SID when evaluating DACL access.
- The CIFS server receives credentials from the CIFS client, including default ACL and user privileges.
  - In contrast, the NFSv4 server receives a more limited set of credentials, and constructs credentials at runtime using its ID mapper. Because of this, the filesystem might see different credentials.

CIFS/SMB Active Directory Integration

The protection system NFSv4 server can be configured to use the Windows Active Directory configuration that is set with the protection system CIFS server.

The system is mapped to use Active Directory if possible. This functionality is disabled by default, but you can enable it using the following command:

```
nfs option set nfs4-idmap-active-directory enabled
```

Default DACL for NFSv4

NFSv4 sets a different default DACL (discretionary access control list) than the default DACL supplied by CIFS.

Only OWNER@, GROUP@ and EVERYONE@ are defined in the default NFSv4 DACL. You can use ACL inheritance to automatically add CIFS-significant ACEs by default if appropriate.

System Default SIDs

Files and directories created by NFSv3, and NFSv4 without ACLs, use the default system domain, sometimes referred to as the default UNIX domain:

- User SIDs in the system domain have format `S-1-22-1-N`, where N is the UID.
- Group SIDs in the system domain have format `S-1-22-2-N`, when N is the GID.

  For example, a user with `UID 1234` would have an owner SID of `S-1-22-1-1234`.

Common identifiers in NFSv4 ACLs and SIDs

The EVERYONE@ identifier and other special identifiers (such as BATCH@, for example) in NFSv4 ACLs use the equivalent CIFS SIDS and are compatible.

The OWNER@ and GROUP@ identifiers have no direct correspondence in CIFS; they appear as the current owner and current owner-group of the file or directory.

NFS Referrals

The referral feature allows an NFSv4 client to access an export (or file system) in one or multiple locations. Locations can be on the same NFS server or on different NFS servers, and use either the same or different path to reach the export.

Because referrals are an NFSv4 feature, they apply only to NFSv4 mounts.

Referrals can be made to any server that uses NFSv4 or later, including the following:

- A protection system running NFS with NFSv4 enabled
- Other servers that support NFSv4 including Linux servers, NAS appliances, and VNX systems.

A referral can use an NFS export point with or without a current underlying path in the DD file system.

NFS exports with referrals can be mounted through NFSv3, but NFSv3 clients will not be redirected since referrals are a NFSv4 feature. This characteristic is useful in scaleout systems to allow exports to be redirected at a file-management level.

Referral Locations

NFSv4 referrals always have one or more locations.

These locations consist of the following:

- A path on a remote NFS server to the referred filesystem.
- One or more server network addresses that allow the client to reach the remote NFS server.

Typically when multiple server addresses are associated with the same location, those addresses are found on the same NFS server.

Referral location names

You can name each referral location within an NFS export. You can use the name to access the referral as well as to modify or delete it.

A referral name can contain a maximum of 80 characters from the following character sets:

- a-z
- A-Z
- 0-9
- "."
- ","
- "_"
- "-"

NOTE: You can include spaces as long as those spaces are embedded within the name. If you use embedded spaces, you must enclose the entire name in double quotes.

Names that begin with "." are reserved for automatic creation by the protection system. You can delete these names but you cannot create or modify them using the command line interface (CLI) or system management services (SMS).

Referrals and Scaleout Systems

NFSv4 referrals and locations can better enable access if you are scaling out your protection systems.

Because your system might or might not already contain a global namespace, the following two scenarios describe how you might use NFSv4 referrals:

- Your system does not contain a global namespace.
  - You can use NFSv4 referrals to build that global namespace. System administrators can build these global namespaces, or you can use smart system manager (SM) element building referrals as necessary.
- Your system already has a global namespace.
  - If your system has a global namespace with MTrees placed in specific nodes, NFS referrals can be created to redirect access to those MTrees to the nodes added to the scaled-out system. You can create these referrals or have them performed automatically within NFS if the necessary SM or file manager (FM) information is available.

NFSv4 and High Availability

With NFSv4, protocol exports (for example, `/data/col1/<mtree>` are mirrored in a High Availability (HA) setup. However, configuration exports such as `/ddvar` are not mirrored.

The `/ddvar` filesystem is unique to each node of an HA pair. As a result, `/ddvar` exports and their associated client access lists are not mirrored to the standby node in an HA environment.

The information in `/ddvar` becomes stale when the active node fails over to the standby node. Any client permissions granted to `/ddvar` on the original active node must be recreated on the newly active node after a failover occurs.

You must also add any additional /ddvar exports and their clients (for example, `/ddvar/core`) that were created on the original active node to the newly active node after a failover occurs.

Finally, any desired `/ddvar` exports must be unmounted from the client and then remounted after a failover occurs.

NFSv4 Global Namespaces

The NFSv4 server provides a virtual directory tree known as a PseudoFS to connect NFS exports into a searchable set of paths.

The use of a PseudoFS distinguishes NFSv4 from NFSv3, which uses the MOUNTD auxiliary protocol.

In most configurations, the change from NFSv3 MOUNTD to NFSv4 global namespace is transparent and handled automatically by the NFSv4 client and server.

NFSv4 global namespaces and NFSv3 submounts

If you use NFSv3 export submounts, the global namespaces characteristic of NFSv4 might prevent submounts from being seen on the NFSv4 mount.

NFSv3 main exports and submount exports

If NFSv3 has a main export and a submount export, these exports might use the same NFSv3 clients yet have different levels of access:

Table 25. NFSv3 main exports and submount exports

Export	Path	Client	Options
Mt1	`/data/col1/mt1`	client1.example.com	ro
Mt1-sub	`/data/col1/mt1/subdir`	client1.example.com	rw

In the previous table, the following applies to NFSv3:

- If client1.example.com mounts /data/col1/mt1, the client gets read-only access.
- If client1.example.com mounts /data/col1/mt1/subdir, the client gets read-write access.

NFSv4 operates in the same manner in regard to highest-level export paths. For NFSv4, client1.example.com navigates the NFSv4 PseudoFS until it reaches the highest-level export path, `/data/col1/mt1`, where it gets read-only access.

However, because the export has been selected, the submount export (Mt1-sub) is not part of the PseudoFS for the client and read-write access is not given.

Best practice

If your system uses NFSv3 exports submounts to give the client read-write access based on the mount path, you must consider this before using NFSv4 with these submount exports.

With NFSv4, each client has an individual PseudoFS.

Table 26. NFSv3 submount exports

Export	Path	Client	Options
Mt1	`/data/col1/mt1`	client1.example.com	ro
Mt1-sub	`/data/col1/mt1/subdir`	client2.example.com	rw

NFSv4 Configuration

The default protection system configuration only enables NFSv3. To use NFSv4, you must first enable the NFSv4 server.

Enabling the NFSv4 Server

Steps

1. Enter `nfs enable version 4` to enable NFSv4:

```
# nfs enable version 4
NFS server version(s) 3:4 enabled.
```

2. (Optional) If you want to disable NFSv3, enter `nfs disable version 3`.

   NOTE: Do not disable NFSv3 on systems integrated with Avamar, or systems that are backup targets for PowerProtect Data Manager.

```
# nfs disable version 3
NFS server version(s) 3 disabled.
NFS server version(s) 4 enabled.
```

Next steps

After the NFSv4 server is enabled, you might need to perform additional NFS configuration tasks specifically for your site. These tasks can include:

- Setting the NFSv4 domain
- Configuring NFSv4 ID mapping
- Configuring ACL (Access Control Lists)

Setting the default server to include NFSv4

About this task

The NFS command option **`default-server-version`** controls which NFS version is enabled when you enter the `nfs enable` command without specifying a version.

Steps

Enter the `nfs option set default-server-version 3:4` command:

```
# nfs option set default-server-version 3:4
NFS option 'default-server-version' set to '3:4'.
```

Updating existing exports

You can update existing exports to change the NFS version used by your protection system.

Steps

Enter the `nfs export modify all` command:

```
# nfs export modify all clients all options version=version number
```

To ensure all existing clients have either version 3, 4, or both, you can modify the NFS version to the appropriate string. The following example shows NFS modified to include versions 3 and 4:

```
#nfs export modify all clients all options version=3:4
```

For more information about the `nfs export` command, see the *DDOS Command Reference Guide* for more information.

Kerberos and NFSv4

Both NFSv4 and NFSv3 use the Kerberos authentication mechanism to secure user credentials.

Kerberos prevents user credentials from being spoofed in NFS packets and protects them from tampering en route to the protection system.

There are distinct types of Kerberos over NFS:

- Kerberos 5 (**`sec=krb5`**)

  Use Kerberos for user credentials.

- Kerberos 5 with integrity (**`sec=krb5i`**)

  Use Kerberos and check the integrity of the NFS payload using an encrypted checksum.

- Kerberos 5 with security (**`sec=krb5p`**)

  Use Kerberos 5 with integrity and encrypt the entire NFS payload.

NOTE: **`krb5i`** and **`krb5p`** can both cause performance degradation due to additional computational overhead on both the NFS client and the protection system.

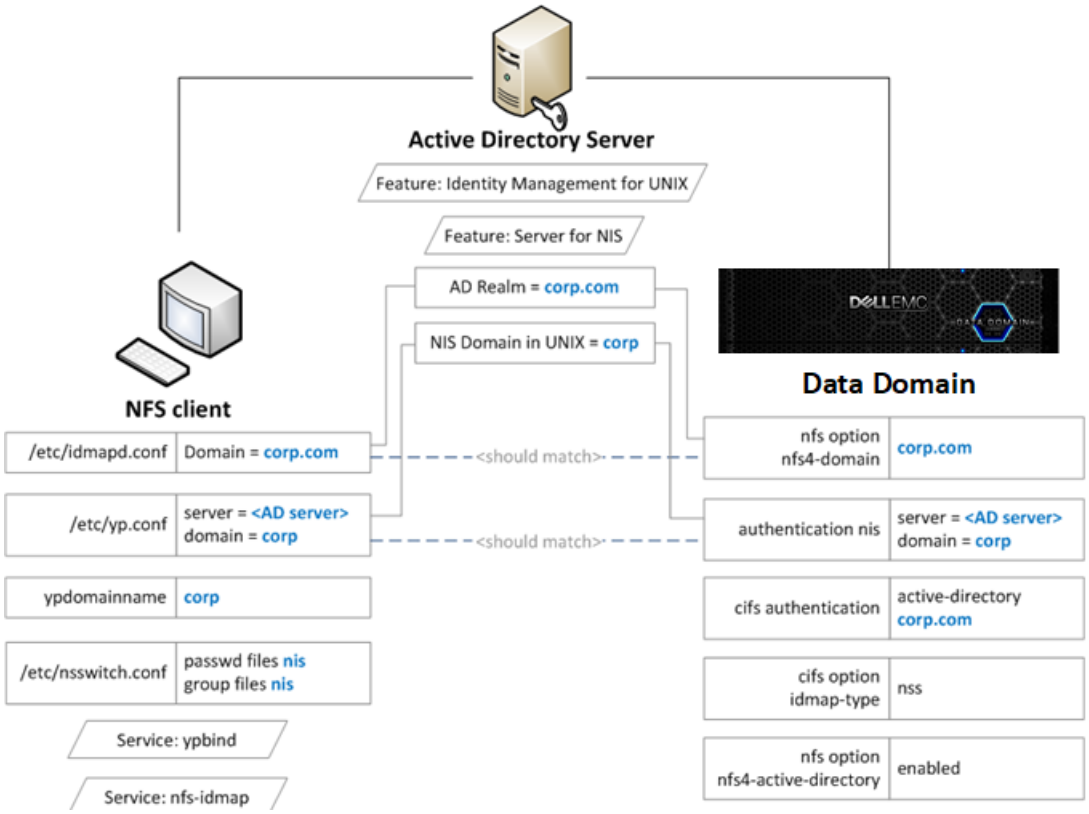

Figure 10. Active Directory Configuration

You employ existing commands that are used for NFSv3 when configuring your system for Kerberos. See the nfsv3 chapter of the *DDOS Command Reference Guide* for more information.

Configuring Kerberos with a Linux-Based KDC

Prerequisites

You should ensure that all your systems can access the Key Distribution Center (KDC).

If the systems cannot reach the KDC, check the domain name system (DNS) settings.

About this task

The following steps allow you to create keytab files for the client and the protection system:

- In Steps 1-3, you create the keytab file for the protection system.
- In Steps 4-5, you create the keytab file for the client.

Steps

1. Create the **nfs/<ddr_dns_name>@<realm>** service principal.

   ```
   kadmin.local: addprinc -randkey nfs/ddr12345.<domain-name>@<domain-name>
   ```

2. Export **nfs/<ddr_dns_name>@<realm>** to a keytab file.

   ```
   kadmin.local: ktadd -k /tmp/ddr.keytab nfs/ddr12345.corp.com@CORP.COM
   ```

3. Copy the keytab file to the protection system at the following location:

   ```
   /ddr/var/krb5.keytab
   ```

4. Create one of the following principals for the client and export that principal to the keytab file:

```
nfs/<client_dns_name>@<REALM>
root/<client_dns_name>@<REALM>
```

5. Copy the keytab file to the client at the following location:

```
/etc/krb5.keytab
```

NOTE: It is recommended that you use an NTP server to keep the time synchronized on all entities.

Configuring the protection System to Use Kerberos Authentication

Steps

1. Configure the KDC and Kerberos realm on the protection system by using the `authentication` command:

```
# authentication kerberos set realm <realm> kdc-type unix kdcs <kdc-server>
```

2. Import the keytab file:

```
# authentication kerberos keytab import
```

3. (Optional) Configure the NIS server by entering the following commands:

```
# authentication nis servers add <server>
# authentication nis domain set <domain-name>
# authentication nis enable
# filesys restart
```

4. (Optional) Make the **nfs4-domain** the same as the Kerberos realm using the `nfs option` command:

```
nfs option set nfs4-domain <kerberos-realm>
```

5. Add a client to an existing export by adding **sec=krb5** to the `nfs export add` command:

```
nfs export add <export-name> clients * options version=4,sec=krb5
```

Configuring Clients

Steps

1. Configure the DNS server and verify that forward and reverse lookups are working.
2. Configure the KDC and Kerberos realm by editing the `/etc/krb5.conf` configuration file.
   You might need to perform this step based on the client operating system you are using.
3. Configure NIS or another external name mapping service.
4. (Optional) Edit the `/etc/idmapd.conf` file to ensure it is the same as the Kerberos realm.
   You might need to perform this step based on the client operating system you are using.
5. Verify the keytab file **/etc/krb5.keytab** contains an entry for the **nfs/** service principal or the **root/** principal.

```
[root@fc22 ~]# klist -k
Keytab name: FILE:/etc/krb5.keytab
KVNO Principal
---- --------------------------------------------------------------------------
3 nfs/fc22.domain-name@domain-name
```

6. Mount the export using the **sec=krb5** option.

```
[root@fc22 ~]# mount ddr12345.<domain-name>:/data/col1/mtree1 /mnt/nfs4 -o
sec=krb5,vers=4
```

Enabling Active Directory

About this task

Configuring Active Directory authentication makes the protection system part of a Windows Active Directory realm. CIFS clients and NFS clients use Kerberos authentication.

Steps

1. Join an active directory realm using the `cifs set` command:

   ```
   # cifs set authentication active-directory <realm>
   ```

   Kerberos is automatically set up on the system, and the required NFS/ service principal is automatically created on the KDC.
2. Configure NIS using the `authentication nis` command:

   ```
   # authentication nis servers add <windows-ad-server>
   # authentication nis domain set <ad-realm>
   # authentication nis enable
   ```

3. Configure CIFS to use NSS for ID mapping by using `cifs` commands:

   ```
   # cifs disable
   # cifs option set idmap-type nss
   # cifs enable
   # filesys restart
   ```

4. Set the **`nfs4-domain`** to be the same as the Active Directory realm:

   ```
   # nfs option set nfs4-domain <ad-realm>
   ```

5. Enable Active Directory for NFSv4 id mapping by using the `nfs` command:

   ```
   # nfs option set nfs4-idmap-active-directory enabled
   ```

Configuring Active Directory

Steps

1. Install the Active Directory Domain Services (AD DS) role on the Windows server.
2. Install the Identity Management for UNIX components.

   ```
   C:\Windows\system32>Dism.exe /online /enable-feature /featurename:adminui /all
   C:\Windows\system32>Dism.exe /online /enable-feature /featurename:nis /all
   ```

3. Verify the NIS domain is configured on the server.

   ```
   C:\Windows\system32>nisadmin
   The following are the settings on localhost

   Push Interval : 1 days
   Logging Mode  : Normal

   NIS Domains
   NIS Domain in AD  Master server  NIS Domain in UNIX
   ----------------  -------------  ----------------
   corp              win-ad-server  corp
   ```

4. Assign AD users and groups UNIX UID/GIDs for the NFSv4 server.
   a. Go to **Server Manager** > **Tools** > **Active Directory**.
   b. Open the **Properties** for an AD user or group.
   c. Under the UNIX Attributes tab, fill in the NIS domain, UID, and Primary GID fields.

Configuring clients on Active Directory

Steps

1. Create a new AD user on the AD server to represent the NFS client's service principal.
2. Create the nfs/ service principal for the NFS client.

```
> ktpass -princ nfs/<client_dns_name>@<REALM> -mapuser nfsuser -pass **** -out
nfsclient.keytab
/crytp rc4-hmac-nt /ptype KRB5_NT_PRINCIPAL
```

3. (Optional) Copy the keytab file to `/etc/krb5.keytab` on the client.

   The need to perform this step depends on which client OS you are using.

19

Metadata on Flash

This chapter includes:

Topics:

Overview of Metadata on Flash (MDoF)

MDoF creates caches for file system metadata using flash technologies. The SSD Cache is a low latency, high input/output operations per second (IOPS) cache to accelerate metadata and data access.

NOTE: The minimum software version required is DDOS 6.0.

Caching the file system metadata on SSDs improves I/O performance for both traditional and random workloads.

For traditional workloads, offloading random access to metadata from HDDs to SSDs allows the hard drives to accommodate streaming write and read requests.

For random workloads, SSD cache provides low latency metadata operations, which allows the HDDs to serve data requests instead of cache requests.

Read cache on SSD improves random read performance by caching frequently accessed data. Writing data to NVRAM combined with low latency metadata operations to drain the NVRAM faster improve random write latency. The absence of cache does not prevent file system operation, it only impacts file system performance.

When the cache tier is first created, a file system restart is only required if the cache tier is being added after the file system is running. For new systems that come with cache tier disks, no file system restart is required if the cache tier is created before enabling the file system for the first time. Additional cache can be added to a live system, without the need to disable and enable the file system.

One specific condition with regard to SSDs is when the number of spare blocks remaining gets close to zero, the SSD enters a read only condition. When a read only condition occurs, DDOS treats the drive as read-only cache and sends an alert.

MDoF is supported on the following systems:

- DD6300
- DD6800
- DD6900
- DD9300
- DD9400
- DD9800
- DD9900
- DD VE instances, including DD3300 systems, in capacity configurations of 16 TB and higher (SSD Cache Tier for DD VE)

SSD cache licensing and capacity

Depending on your system model, the SSD cache feature will either be enabled by default with no need for a license, or it will require an ELMS license to enable

The following table describes the various SSD capacity and licensing requirements for the supported systems:

Table 27. SSD capacity and licensing requirements

Model	Memory	Number of SSDs	SSD capacity	License required	Enabled by default
DD6300	48 GB (Base)	1	800 GB	Y	N
	96 GB (Expanded)	2	1600 GB	Y	N
DD6800	192 GB (Base)	2	1600 GB	Y	N
	192 GB (Expanded)	4	3200 GB	Y	N
DD6900	288 GB	2	3840 GB	N	Y
DD9300	192 GB (Base)	5	4000 GB	Y	N
	384 GB (Expanded)	8	6400 GB	Y	N
DD9400	576 GB	5	19200 GB	N	Y
DD9800	256 GB (Base)	8	6400 GB	Y	N
	768 GB (Expanded)	15	12000 GB	Y	N
DD9900	1152 GB	10	38400 GB	N	Y

SSD Cache Tier for DD VE

DD VE instances and DD3300 systems do not require a license for the SSD Cache Tier. The maximum supported SSD capacity is 1% of the Active Tier capacity.

The following table describes the various SSD capacity licenses and the SSD capacities for the given system:

Table 28. DD VE and DD3300 SSD capacity

Capacity configuration	Maximum SSD capacity
DD VE 16 TB	160 GB
DD VE 32 TB	320 GB
DD VE 48 TB	480 GB
DD VE 64 TB	640 GB
DD VE 96 TB	960 GB
DD3300 8 TB	160 GB
DD3300 16 TB	160 GB
DD3300 32 TB	320 GB

SSD cache tier

The SSD cache tier provides the SSD cache storage for the file system. The file system draws the required storage from the SSD cache tier without active intervention from the user.

SSD cache tier - system management

Be aware of the following considerations for SSD cache:

- When SSDs are deployed within a controller, those SSDs are treated as internal root drives. They display as enclosure 1 in the output of the `storage show all` command.
- Manage individual SSDs with the `disk` command the same way HDDs are managed.
- Run the `storage add` command to add an individual SSD or SSD enclosure to the SSD cache tier.
- The SSD cache tier space does not need to be managed. The file system draws the required storage from the SSD cache tier and shares it among its clients.
- The `filesys create` command creates an SSD volume if SSDs are available in the system.

  NOTE: If SSDs are added to the system later, the system should automatically create the SSD volume and notify the file system. SSD Cache Manager notifies its registered clients so they can create their cache objects.
- If the SSD volume contains only one active drive, the last drive to go offline will come back online if the active drive is removed from the system.

The next section describes how to manage the SSD cache tier from DD System Manager, and with the DDOS CLI.

Managing the SSD cache tier

Storage configuration features allow you to add and remove storage from the SSD cache tier.

Steps

1. Select **Hardware** > **Storage** > **Overview**.
2. Expand the **Cache Tier** dialog.
3. Click **Configure**.

   NOTE: The licensed capacity bar shows the portion of licensed capacity (used and remaining) for the installed enclosures.
4. Select the checkbox for the Shelf to be added.
5. Click the **Add to Tier** button.
6. Click **OK** to add the storage.

   NOTE: To remove an added shelf, select it in the Tier Configuration list, click **Remove from Configuration**, and click **OK**.

   CLI Equivalent

   When the cache tier SSDs are installed in the head unit:

   a. Add the SSDs to the cache tier.

```
#  storage add disks 1.13,1.14 tier cache
Checking storage requirements...done
Adding disk 1.13 to the cache tier...done

Updating system information...done

Disk 1.13 successfully added to the cache tier.


Checking storage requirements...
done
Adding disk 1.14 to the cache tier...done

Updating system information...done

Disk 1.14 successfully added to the cache tier.
```

   b. Verify the state of the newly added SSDs.

```
# disk show state
Enclosure   Disk
             1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
---------   ------------------------------------------------
1            .  .  .  .  s  .  .  s  s  s  s  s  v  v
2            U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
3            U  U  U  U  U  U  U  U  U  U  U  U  U  U  U
---------   ------------------------------------------------
```

```
Legend   State              Count
------   ---------------    -----
.        In Use Disks       6
s        Spare Disks        6
v        Available Disks    2
U        Unknown Disks      30
------   ---------------    -----
Total 44 disks
```

When the cache tier SSDs are installed in an external shelf:

a. Verify the system recognizes the SSD shelf. In the example below, the SSD shelf is enclosure 2.

```
# disk show state
Enclosure   Disk
Row(disk-id) 1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
---------    ----------------------------------------------
1            .  .  .  .
2            U  U  U  U  U  U  U  U  -  -  -  -  -  -  -
3            .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
4            .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
5            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
7            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
8            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
9            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
10          |--------|--------|--------|--------|
            | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
   E(49-60) |v  v  v |v  v  v |v  v  v |v  v  v |
   D(37-48) |v  v  v |v  v  v |v  v  v |v  v  v |
   C(25-36) |v  v  v |v  v  v |v  v  v |v  v  v |
   B(13-24) |v  v  v |v  v  v |v  v  v |v  v  v |
   A( 1-12) |v  v  v |v  v  v |v  v  v |v  v  v |
            |--------|--------|--------|--------|
11           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
12           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
13           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
---------    ----------------------------------------------

Legend   State                Count
------   ------------------   -----
.        In Use Disks         32
v        Available Disks      182
U        Unknown Disks        8
-        Not Installed Disks  7
------   ------------------   -----
Total 222 disks
```

b. Identify the shelf ID of the SSD shelf. SSDs will display as `SAS-SSD` or `SATA-SSD` in the `Type` column.

```
#  disk show hardware
```

Disk (enc/disk)	Slot	Manufacturer/Model	Firmware	Serial No.	Capacity	Type
1.1	0	TG32C10400GA3EMC_____118000371	PRO6E344	FG009826	372.61 GiB	SATA-SSD
1.2	1	TG32C10400GA3EMC_____118000371	PRO6E344	FG0097VL	372.61 GiB	SATA-SSD
1.3	2	TG32C10400GA3EMC_____118000371	PRO6E344	FG009881	372.61 GiB	SATA-SSD
1.4	3	TG32C10400GA3EMC_____118000371	PRO6E344	FG00988X	372.61 GiB	SATA-SSD
2.1	0	HITACHI HUSMR148_CLAR800	C29C	07V4P2AA	745.22 GiB	SAS-SSD
2.2	1	HITACHI HUSMR148_CLAR800	C29C	07V4P3LA	745.22 GiB	SAS-SSD
2.3	2	HITACHI HUSMR148_CLAR800	C29C	07V4P2XA	745.22 GiB	SAS-SSD
2.4	3	HITACHI HUSMR148_CLAR800	C29C	07V4TW4A	745.22 GiB	SAS-SSD
2.5	4	HITACHI HUSMR148_CLAR800	C29C	07V4ULYA	745.22 GiB	SAS-SSD
2.6	5	HITACHI HUSMR148_CLAR800	C29C	07V4P0BA	745.22 GiB	SAS-SSD
2.7	6	HITACHI HUSMR148_CLAR800	C29C	07V4UVBA	745.22 GiB	SAS-SSD
2.8	7	HITACHI HUSMR148_CLAR800	C29C	07V4UTNA	745.22 GiB	SAS-SSD

Figure 11.

c. Add the SSD shelf to the cache tier

```
#  storage add enclosure 2 tier cache

Checking storage requirements...done
Adding enclosure 2 to the cache tier...Enclosure 2 successfully added to the cache
tier.
```

```
Updating system information...done

Successfully added: 2  done
```

d. Verify the state of the newly added SSDs.

```
# disk show state
Enclosure   Disk
Row(disk-id) 1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
---------    ---------------------------------------------
1            .  .  .  .
2            .  .  .  .  .  .  .  .  -  -  -  -  -  -  -
3            .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
4            .  .  .  .  .  .  .  .  .  .  .  .  .  .  v
5            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
6            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
7            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
8            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
9            v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
10           |--------|--------|--------|--------|
             | Pack 1 | Pack 2 | Pack 3 | Pack 4 |
   E(49-60)  |v  v  v |v  v  v |v  v  v |v  v  v |
   D(37-48)  |v  v  v |v  v  v |v  v  v |v  v  v |
   C(25-36)  |v  v  v |v  v  v |v  v  v |v  v  v |
   B(13-24)  |v  v  v |v  v  v |v  v  v |v  v  v |
   A( 1-12)  |v  v  v |v  v  v |v  v  v |v  v  v |
             |--------|--------|--------|--------|
11           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
12           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
13           v  v  v  v  v  v  v  v  v  v  v  v  v  v  v
---------    ---------------------------------------------

Legend   State               Count
------   ------------------  -----
.        In Use Disks        32
v        Available Disks     182
U        Unknown Disks       8
-        Not Installed Disks 7
------   ------------------ -----
Total 222 disks
```

To remove a controller-mounted SSD from the cache tier:

```
# storage remove disk 1.13

Removing disk 1.13...done

Updating system information...done

Disk 1.13 successfully removed.
```

To remove an SSD shelf from the system:

```
# storage remove enclosure 2

Removing enclosure 2...Enclosure 2 successfully removed.

Updating system information...done

Successfuly removed: 2 done
```

SSD alerts

There are three alerts specific to the SSD cache tier.

The SSD cahce tier alerts are:

- Licensing

If the file system is enabled and less physical cache capacity present than what the license permits is configured, an alert is generated with the current SSD capacity present, and the capacity license. This alert is classified as a warning alert. The absence of cache does not prevent file system operation, it only impacts file system performance. Additional cache can be added to a live system, without the need to disable and enable the file system.

- Read only condition

  When the number of spare blocks remaining gets close to zero, the SSD enters a read only condition. When a read only condition occurs, DDOS treats the drive as read-only cache.

  Alert `EVT-STORAGE-00001` displays when the SSD is in a read-only state and should be replaced.

- SSD end of life

  When an SSD reaches the end of its lifespan, the system generates a hardware failure alert identifying the location of the SSD within the SSD shelf. This alert is classified as a critical alert.

  Alert `EVT-STORAGE-00016` displays when the EOL counter reaches 98. The drive is failed proactively when the EOL counter reaches 99.

20

SCSI Target

This chapter includes:

Topics:

SCSI Target overview

SCSI (Small Computer System Interface) Target is a unified management daemon for all SCSI services and transports. SCSI Target supports DD VTL (Virtual Tape Library), DD Boost over FC (Fibre Channel), and vDisk/ProtectPoint Block Services, as well as anything that has a target LUN (logical unit number) on a DD system.

SCSI Target Services and Transports

The SCSI Target daemon starts when FC ports are present or DD VTL is licensed. It provides unified management for all SCSI Target *services* and *transports*.

- A *service* is anything that has a target LUN on a DD system that uses SCSI Target commands, such as DD VTL (tape drives and changers), DD Boost over FC (processor devices), or vDisk (Virtual Disk Device).
- A *transport* enables *devices* to become visible to *initiators*.
- An *initiator* is a backup client that connects to a system to read and write data using the FC protocol. A specific initiator can support DD Boost over FC, vDisk, or DD VTL, but not all three.
- *Devices* are visible on a SAN (storage area network) through physical ports. Host initiators communicate with the DD system through the SAN.
- *Access groups* manage access between devices and initiators.
- An *endpoint* is the logical target on a DD system to which an initiator connects. You can disable, enable, and rename endpoints. To delete endpoints, the associated transport hardware must no longer exist. Endpoints are automatically discovered and created when a new transport connection occurs. Endpoints have the following attributes: port topology, FCP2-RETRY status, WWPN, and WWNN.
- *NPIV* (N_port ID Virtualization) is an FC feature that lets multiple endpoints share a single physical port. NPIV eases hardware requirements and provides failover capabilities.
- In DDOS, users can specify the sequence of secondary system addresses for failover. For example, if the system specifies 0a, 0b, 1a, 1b and the user specifies 1b, 1a, 0a, 0b , the user-specified sequence is used for failover. The `scsitarget endpoint show detailed` command displays the user-specified sequence.

Only one initiator should be present per access group. Each access group is assigned a type (DD VTL, vDisk/ProtectPoint Block Services, or DD Boost over FC).

SCSI Target Architectures - Supported and Unsupported

SCSI Target supports the following architectures:

- **DD VTL plus DD Boost over FC from different initiators:** Two different initiators (on the same or different clients) may access a DD system using DD VTL and DD Boost over FC, through the same or different DD system target endpoints.
- **DD VTL plus DD Boost over FC from one initiator to two different DD systems:** A single initiator may access two different DD systems using any service.

SCSI Target does not support the following architecture:

- **DD VTL plus DD Boost over FC from one initiator to the same DD system:** A single initiator may not access the same DD system through different services.

Thin Protocol

The thin protocol is a lightweight daemon for VDisk and DD VTL that responds to SCSI commands when the primary protocol can't. For Fibre Channel environments with multiple protocols, thin protocol:

- Prevents initiator hangs
- Prevents unnecessary initiator aborts
- Prevents initiator devices from disappearing
- Supports a standby mode
- Supports fast and early discoverable devices
- Enhances protocol HA behavior
- Doesn't require fast registry access

For More Information about DD Boost and the scscitarget Command (CLI)

For more information about using DD Boost through the DD System Manager, see the related chapter in this book. For other types of information about DD Boost, see the *DD Boost for OpenStorage Administration Guide*.

This chapter focuses on using SCSI Target through the DD System Manager. After you have become familiar with basic tasks, the `scscitarget` command in the *DDOS Command Reference Guide* provides more advanced management tasks.

When there is heavy DD VTL traffic, avoid running the `scsitarget group use` command, which switches the in-use endpoint lists for one or more SCSI Target or vdisk devices in a group between primary and secondary endpoint lists.

Related concepts

About DD Boost

Fibre Channel view

The Fibre Channel view displays the current status of whether Fibre Channel and/or NPIV is enabled. It also displays two tabs: Resources and Access Groups. Resources include ports, endpoints, and initiators. An access group holds a collection of initiator WWPNs (worldwide port names) or aliases and the drives and changers they are allowed to access.

Enable N_Port ID Virtualization on a Data Domain system

Prerequisites

The following conditions must be met:

- The Data Domain system must be running a supported version of Data Domain OS.
- All ports must be connected to 4 Gb, 8 Gb, and 16 Gb Fibre Channel HBA and SLIC.
- The Data Domain system ID must be valid, that is, it must not be 0.

In addition, port topologies and port names are reviewed and may prevent NPIV from being enabled:

- NPIV is allowed if the topology for *all* ports is loop-preferred.
- NPIV is allowed if the topology for *some* of the ports is loop-preferred. However, NPIV must be disabled for ports that are loop-only, or you must reconfigure the topology to loop-preferred for proper functionality.
- NPIV is *not* allowed if *none* of the ports have a topology of loop-preferred.
- If port names are present in access groups, the port names are replaced with their associated endpoint names.

Steps

1. Select **Hardware** > **Fibre Channel**.
2. Next to **NPIV: Disabled**, select **Enable**.
3. In the **Enable NPIV** dialog box, you are warned that all Fibre Channel ports must be disabled before NPIV can be enabled. If you are sure that you want to continue, select **Yes**.

Disabling NPIV

Before you can disable NPIV, you must not have any ports with multiple endpoints.

About this task

NOTE: NPIV is required for HA configuration. It is enabled by default and cannot be disabled.

Steps

1. Select **Hardware** > **Fibre Channel**.
2. Next to NPIV: Enabled, select **Disable**.
3. In the Disable NPIV dialog, review any messages about correcting the configuration, and when ready, select **OK**.

Resources tab

The **Hardware** > **Fibre Channel** > **Resources** tab displays information about ports, endpoints, and initiators.

Table 29. Ports

Item	Description
System Address	System address for port
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node
Enabled	Port operational status; either Enabled or Disabled.
NPIV	NPIV status; either Enabled or Disabled.
Link Status	Link status: either Online or Offline; that is, whether or not the port is up and capable of handling traffic.
Operation Status	Operation status: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this port.

Table 30. Endpoints

Item	Description
Name	Name of endpoint.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node
System Address	System address of endpoint.
Enabled	Port operational state; either Enabled or Disabled.
Link Status	Either Online or Offline; that is, whether or not the port is up and capable of handling traffic.

Table 31. Initiators

Item	Description
Name	Name of initiator.
Service	Service support by the initiator, which is either DD VTL, DD Boost, or vDisk.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
Vendor Name	Initiator's model.
Online Endpoints	Endpoints seen by this initiator. Displays `none` or `offline` if the initiator is not available.

Related concepts

Working with endpoints
Working with initiators

Configuring a port

Ports are discovered, and a single endpoint is automatically created for each port, at startup.

About this task

The properties of the base port depend on whether NPIV is enabled:

- In non-NPIV mode, ports use the same properties as the endpoint, that is, the WWPN for the base port and the endpoint are the same.
- In NPIV mode, the base port properties are derived from default values, that is, a new WWPN is generated for the base port and is preserved to allow consistent switching between NPIV modes. Also, NPIV mode provides the ability to support multiple endpoints per port.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Under **Ports**, select an port, and then select **Modify** (pencil).
3. In the Configure Port dialog, select whether to automatically enable or disable NPIV for this port.
4. For Topology, select Loop Preferred, Loop Only, Point to Point, or Default.
5. For Speed, select 1, 2, 4, 8, or 16 Gbps, or auto.
6. Select **OK**.

Enabling a port

Ports must be enabled before they can be used.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Ports** > **Enable**. If all ports are already enabled, a message to that effect is displayed.
3. In the Enable Ports dialog, select one or more ports from the list, and select **Next**.
4. After the confirmation, select **Next** to complete the task.

Disabling a port

You can simply disable a port (or ports), or you can chose to failover all endpoints on the port (or ports) to another port.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Ports** > **Disable**.
3. In the Disable Ports dialog, select one or more ports from the list, and select **Next**.
4. In the confirmation dialog, you can continue with simply disabling the port, or you can chose to failover all endpoints on the ports to another port.

Adding an endpoint

An endpoint is a virtual object that is mapped to a underlying virtual port. In non-NPIV mode (not available on HA configuration), only a single endpoint is allowed per physical port, and the base port is used to configure that endpoint to the fabric. When NPIV is enabled, multiple endpoints are allowed per physical port, each using a virtual (NPIV) port, and endpoint failover/failback is enabled.

About this task

NOTE: Non-NPIV mode is not available on HA configurations. NPIV is enabled by default and cannot be disabled.

NOTE: In NPIV mode, endpoints:

- have a primary system address.
- may have zero or more secondary system addresses.
- are all candidates for failover to an alternate system address on failure of a port; however, failover to a marginal port is not supported.
- may be failed back to use their primary port when the port comes back up online.

NOTE: When using NPIV, it is recommended that you use only one protocol (that is, DD VTL Fibre Channel, DD Boost-over-Fibre Channel, or vDisk Fibre Channel) per endpoint. For failover configurations, secondary endpoints should also be configured to have the same protocol as the primary.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Under **Endpoints**, select **Add** (+ sign).
3. In the Add Endpoint dialog, enter a Name for the endpoint (from 1 to 128 characters). The field cannot be empty or be the word "all," and cannot contain the characters asterisk (*), question mark (?), front or back slashes (/, \), or right or left parentheses [(,)].
4. For Endpoint Status, select Enabled or Disabled.
5. If NPIV is enabled, for Primary system address, select from the drop-down list. The primary system address must be different from any secondary system address.
6. If NPIV is enabled, for Fails over to secondary system addresses, check the appropriate box next to the secondary system address.
7. Select **OK**.

Configuring an endpoint

After you have added an endpoint, you can modify it using the Configure Endpoint dialog.

About this task

NOTE: When using NPIV, it is recommended that you use only one protocol (that is, DD VTL Fibre Channel, DD Boost-over-Fibre Channel, or vDisk Fibre Channel) per endpoint. For failover configurations, secondary endpoints should also be configured to have the same protocol as the primary.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Under **Endpoints**, select an endpoint, and then select **Modify** (pencil).
3. In the Configure Endpoint dialog, enter a Name for the endpoint (from 1 to 128 characters). The field cannot be empty or be the word "all," and cannot contain the characters asterisk (*), question mark (?), front or back slashes (/, \), or right or left parentheses [(,)].
4. For Endpoint Status, select Enabled or Disabled.
5. For Primary system address, select from the drop-down list. The primary system address must be different from any secondary system address.
6. For Fails over to secondary system addresses, check the appropriate box next to the secondary system address.
7. Select **OK**.

Modifying an endpoint's system address

You can modify the active system address for a SCSI Target endpoint using the `scsitarget endpoint modify` command option. This is useful if the endpoint is associated with a system address that no longer exists, for example after a controller upgrade or when a controller HBA (host bus adapter) has been moved. When the system address for an endpoint is modified, all

properties of the endpoint, including WWPN and WWNN (worldwide port and node names, respectively), if any, are preserved and are used with the new system address.

About this task

In the following example, endpoint ep-1 was assigned to system address 5a, but this system address is no longer valid. A new controller HBA was added at system address 10a. The SCSI Target subsystem automatically created a new endpoint, ep-new, for the newly discovered system address. Because only a single endpoint can be associated with a given system address, ep-new must be deleted, and then ep-1 must be assigned to system address 10a.

NOTE: It may take some time for the modified endpoint to come online, depending on the SAN environment, since the WWPN and WWNN have moved to a different system address. You may also need to update SAN zoning to reflect the new configuration.

Steps

1. Show all endpoints to verify the endpoints to be changed:
   ```
   # scsitarget endpoint show list
   ```
2. Disable all endpoints:
   ```
   # scsitarget endpoint disable all
   ```
3. Delete the new, unnecessary endpoint, ep-new:
   ```
   # scsitarget endpoint del ep-new
   ```
4. Modify the endpoint you want to use, ep-1, by assigning it the new system address 10a:
   ```
   # scsitarget endpoint modify ep-1 system-address 10a
   ```
5. Enable all endpoints:
   ```
   # scsitarget endpoint enable all
   ```

Enabling an endpoint

Enabling an endpoint enables the port only if it is currently disabled, that is, you are in non-NPIV mode.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Endpoints** > **Enable**. If all endpoints are already enabled, a message to that effect is displayed.
3. In the Enable Endpoints dialog, select one or more endpoints from the list, and select **Next**.
4. After the confirmation, select **Next** to complete the task.

Disabling an endpoint

Disabling an endpoint does not disable the associated port, unless all endpoints using the port are disabled, that is, you are in non- NPIV mode.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Endpoints** > **Disable**.
3. In the Disable Endpoints dialog, select one or more endpoints from the list, and select **Next**. If an endpoint is in use, you are warned that disabling it might disrupt the system.
4. Select **Next** to complete the task.

Deleting an endpoint

You may want to delete an endpoint if the underlying hardware is no longer available. However, if the underlying hardware is still present, or becomes available, a new endpoint for the hardware is discovered automatically and configured based on default values.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Endpoints** > **Delete**.
3. In the Delete Endpoints dialog, select one or more endpoints from the list, and select **Next**. If an endpoint is in use, you are warned that deleting it might disrupt the system.
4. Select **Next** to complete the task.

Adding an initiator

Add initiators to provide backup clients to connect to the system to read and write data using the FC (Fibre Channel) protocol. A specific initiator can support DD Boost over FC, or DD VTL, but not both. A maximum of 1024 initiators can be configured for a DD system.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Under Initiators, select Add (+ sign)
3. In the Add Initiator dialog, enter the port's unique WWPN in the specified format.
4. Enter a Name for the initiator.
5. Select the Address Method: **Auto** is used for standard addressing, and **VSA** (Volume Set Addressing) is used primarily for addressing virtual buses, targets, and LUNs.
6. Select **OK**.

CLI Equivalent

```
# scsitarget group add My_Group initiator My_Initiator
```

Modifying or deleting an initiator

Before you can delete an initiator, it must be offline and not attached to any group. Otherwise, you will get an error message, and the initiator will not be deleted. You must delete all initiators in an access group before you can delete the access group. If an initiator remains visible, it may be automatically rediscovered.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Under Initiators, select one of the initiators. If you want to delete it, select Delete (X). If you want to modify it, select Modify (pencil) to display the Modify Initiator dialog.
3. Change the initiator's Name and/or Address Method [**Auto** is used for standard addressing, and **VSA** (Volume Set Addressing) is used primarily for addressing virtual buses, targets, and LUNs.]
4. Select **OK**.

Recommendation to Set Initiator Aliases - CLI only

It is strongly recommended that Initiator aliases be set to reduce confusion and human error during the configuration process.

If this initiator exists, use the `scsitarget initiator rename` command to change its alias

```
#  scsitarget initiator show list
Initiator     System Address            Group   Service
-----------   ----------------------    -----   -------
initiator-1   10:00:00:10:9b:64:d9:2d     n/a       n/a
```

```
-----------   ----------------------   -----   -------

# scsitarget initiator rename initiator-1 NewAliasName
Initiator 'initiator-1' successfully renamed.

# scsitarget initiator show list
Initiator      System Address            Group   Service
------------   ----------------------   -----   -------
NewAliasName   10:00:00:10:9b:64:d9:2d     n/a       n/a
------------   ----------------------   -----   -------
```

If this initiator does not exist, run the `scsitarget initiator add` command to add it and give it an alias.

```
#  scsitarget initiator show list
No matching initiators were found.

# scsitarget initiator add NewAliasName system-address 21:00:00:e0:8b:9d:0b:e8
Initiator 'NewAliasName' successfully added.

# scsitarget initiator show list
Initiator      System Address            Group   Service
------------   ----------------------   -----   -------
NewAliasName   21:00:00:e0:8b:9d:0b:e8     n/a       n/a
------------   ----------------------   -----   -------
```

Setting a hard address (loop ID)

Some backup software requires that all private-loop targets have a hard address (loop ID) that does not conflict with another node. The range for a loop ID is from 0 to 125.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Set Loop ID**.
3. In the Set Loop ID dialog, enter the loop ID (from 0 to 125), and select **OK**.

Setting failover options

You can set options for automatic failover and failback when NPIV is enabled.

About this task

Here is the expected behavior for Fibre Channel port failover, by application:

- DD Boost-over-Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.
- DD VTL Fibre Channel operation is expected to be interrupted when the DD VTL Fibre Channel endpoints failover. You may need to perform discovery (that is, operating system discovery and configuration of DD VTL devices) on the initiators using the affected Fibre Channel endpoint. You should expect to re-start active backup and restore operations.
- vDisk Fibre Channel operation is expected to continue without user intervention when the Fibre Channel endpoints failover.

Automatic failback is not guaranteed if all ports are disabled and then subsequently enabled (which could be triggered by the administrator), as the order in which ports get enabled is unspecified.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources**.
2. Select **More Tasks** > **Set Failover Options**.
3. In the Set Failover Options dialog, enter the Failover and Failback Delay (in seconds) and whether to enable Automatic Failback, and select **OK**.

Access Groups tab

The **Hardware** > **Fibre Channel** > **Access Groups** tab provides information about DD Boost and DD VTL access groups. Selecting the link to *View DD Boost Groups* or *View VTL Groups* takes you to the DD Boost or DD VTL pages.

Table 32. Access Groups

Item	Description
Group Name	Name of access group.
Service	Service for this access group: either DD Boost or DD VTL.
Endpoints	Endpoints associated with this access group.
Initiators	Initiators associated with this access group.
Number of Devices	Number of devices associated with this access group.

Related concepts

Working with access groups

Related tasks

Verifying connectivity and creating access groups

Port monitoring

Port monitoring detects an FC port at system startup and raises an alert if the port is enabled and offline.

To clear the alert, disable an unused port using the `scsitarget port` commands.

21

Working with DD Boost

This chapter includes:

Topics:

About DD Boost

DD Boost provides advanced integration with backup and enterprise applications for increased performance and ease of use. DD Boost distributes parts of the deduplication process to the backup server or application clients, enabling client-side deduplication for faster, more efficient backup and recovery.

DD Boost is an optional product that requires a separate license to operate on the protection system. You can purchase a DD Boost software license key directly from Dell.

NOTE: A special license, BLOCK-SERVICES-PROTECTPOINT, is available to enable clients using ProtectPoint block services to have DD Boost functionality without a DD Boost license. If DD Boost is enabled for ProtectPoint clients only—that is, if only the BLOCK-SERVICES-PROTECTPOINT license is installed—the license status indicates that DD Boost is enabled for ProtectPoint only.

There are two components to DD Boost: one component that runs on the backup server and another that runs on the protection system.

- In the context of the NetWorker backup application, Avamar backup application and other DDBoost partner backup applications, the component that runs on the backup server (DD Boost libraries) is integrated into the particular backup application.
- In the context of Veritas backup applications (NetBackup and Backup Exec) and the Oracle RMAN plug-in, you need to download an appropriate version of the DD Boost plugin that is installed on each media server. The DD Boost plugin includes the DD Boost libraries for integrating with the DD Boost server running on the protection system.

The backup application (for example, Avamar, NetWorker, NetBackup, or Backup Exec) sets policies that control when backups and duplications occur. Administrators manage backup, duplication, and restores from a single console and can use all of the features of DD Boost, including WAN-efficient replicator software. The application manages all files (collections of data) in the catalog, even those created by the protection system.

In the protection system, storage units that you create are exposed to backup applications that use the DD Boost protocol. For Veritas applications, storage units are viewed as disk pools. For Networker, storage units are viewed as logical storage units (LSUs). A storage unit is an MTree; therefore, it supports MTree quota settings. (Do not create an MTree in place of a storage unit.)

This chapter does not contain installation instructions; refer to the documentation for the product you want to install. For example, for information about setting up DD Boost with Veritas backup applications (NetBackup and Backup Exec), see the *DD Boost for OpenStorage Administration Guide*. For information on setting up DD Boost with any other application, see the application-specific documentation.

Additional information about configuring and managing DD Boost on the protection system can also be found in the *DD Boost for OpenStorage Administration Guide* (for NetBackup and Backup Exec) and the *DD Boost for Partner Integration Administration Guide* (for other backup applications).

Managing DD Boost with DD System Manager

Access the DD Boost view in DD System Manager.

Prerequisites

NFSv3 must be enabled to use DD Boost.

Steps

1. Select **Data Management** > **File System**. Verify that the file system is enabled and running by checking its state.
2. Select **Protocols** > **DD Boost**.

   If you go to the DD Boost page without a license, the Status states that DD Boost is not licensed. Click **Add License** and enter a valid license in the Add License Key dialog box.

   NOTE: A special license, BLOCK-SERVICES-PROTECTPOINT, is available to enable clients using ProtectPoint block services to have DD Boost functionality without a DD Boost license. If DD Boost is enabled for ProtectPoint clients only—that is, if only the BLOCK-SERVICES-PROTECTPOINT license is installed—the license status indicates that DD Boost is enabled for ProtectPoint only.

   Use the DD Boost tabs—Settings, Active Connections, IP Network, Fibre Channel, and Storage Units—to manage DD Boost.

Specifying DD Boost user names

A DD Boost user is also a DDOS user. Specify a DD Boost user either by selecting an existing DDOS user name or by creating a new DDOS user name and making that name a DD Boost user.

About this task

Backup applications use the DD Boost user name and password to connect to the protection system. You must configure these credentials on each backup server that connects to this system. The system supports multiple DD Boost users. For complete information about setting up DD Boost with Veritas NetBackup and Backup Exec, see the *DD Boost for OpenStorage Administration Guide*. For information on setting up DD Boost with other applications, see the *DD Boost for Partner Integration Administration Guide* and the application-specific documentation.

Steps

1. Select **Protocols** > **DD Boost**.
2. Select **Add** (+) above the Users with DD Boost Access list.

   The Add User dialog appears.

3. To select an existing user, select the user name in the drop-down list.

   If possible, select a user name with management role privileges set to *none*.

4. To create and select a new user, select **Create a new Local User** and do the following:
   a. Enter the new user name in the User field.

      The user must be configured in the backup application to connect to the protection system.

   b. Enter the password twice in the appropriate fields.
5. Click **Add**.

Changing DD Boost user passwords

Change a DD Boost user password.

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.

2. Select a user in the Users with DD Boost Access list.
3. Click the **Edit** button (pencil icon) above the DD Boost user list.

   The Change Password dialog appears.

4. Enter the password twice in the appropriate boxes.
5. Click **Change**.

Troubleshooting DD Boost user access issues

DD Boost user is locked out

The most common reason a user becomes locked out of the system is that the password expired. Passwords must be changed at intervals specified by the system administrator (90 days by default). Refer to the KB article *Data Domain: DDBoost user shows locked status*, available from the Online Support website for information on resolving and preventing this issue.

Removing a DD Boost user name

Remove a user from the DD Boost access list.

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.
2. Select the user in the Users with DD Boost Access list that needs to be removed.
3. Click **Remove** (X) above the DD Boost user list.

   The Remove User dialog appears.

4. Click **Remove**.

   After removal, the user remains in the DDOS access list.

Related concepts

System access management

Enabling DD Boost

Use the DD Boost Settings tab to enable DD Boost and to select or add a DD Boost user.

Steps

1. Select **Protocols** > **DD Boost**.
2. Click **Enable** in the DD Boost Status area.

   The Enable DD Boost dialog box is displayed.

3. Select an existing user name from the menu, or add a new user by supplying the name, password, and role.

Configuring Kerberos

You can configure Kerberos by using the DD Boost Settings tab.

Steps

1. Select **Protocols > DD Boost > Settings**.
2. Click **Configure** in the Kerberos Mode status area.
   The Authentication tab under **Administration > Access** is displayed.

NOTE: You can also enable Kerberos by going directly to Authentication under **Administration > Access** in System Manager.

3. Under Active Directory/Kerberos Authentication, click **Configure**.
   The Active Directory/Kerberos Authentication dialog box is displayed.

   Choose the type of Kerberos Key Distribution Center (KDC) you want to use:

   - **Disabled**
     NOTE: If you select **Disabled**, NFS clients do not use Kerberos authentication. CIFS clients use Workgroup authentication.
   - **Windows/Active Directory**
     NOTE: Enter the Realm Name, Under Name, and Password for Active Directory authentication.
   - **Unix**
     a. Enter the Realm Name, the IP Address/Host Names of one to three KDC servers.
     b. Upload the keytab file from one of the KDC servers.

Disabling DD Boost

Disabling DD Boost drops all active connections to the backup server. When you disable or destroy DD Boost, the DD Boost FC service is also disabled.

Prerequisites

Ensure there are no jobs running from your backup application before disabling.

About this task

NOTE: File replication started by DD Boost between two restore operations is not canceled.

Steps

1. Select **Protocols** > **DD Boost**.
2. Click **Disable** in the DD Boost Status area.
3. Click **OK** in the Disable DD Boost confirmation dialog box.

Viewing DD Boost storage units

Access the Storage Units tab to view and manage DD Boost storage units.

The DD Boost Storage Unit tab:

- Lists the storage units and provides the following information for each storage unit:

Table 33. Storage unit information

Item	Description
Storage Unit	The name of the storage unit.
User	The DD Boost user owning the storage unit.
Quota Hard Limit	The percentage of hard limit quota used.
Last 24 hr Pre-Comp	The amount of raw data from the backup application that has been written in the last 24 hours.
Last 24 hr Post-Comp	The amount of storage used after compression in the last 24 hours.
Last 24 hr Comp Ratio	The compression ratio for the last 24 hours.
Weekly Avg Post-Comp	The average amount of compressed storage used in the last five weeks.
Last Week Post-Comp	The average amount of compressed storage used in the last seven days.
Weekly Avg Comp Ratio	The average compression ratio for the last five weeks.

Table 33. Storage unit information (continued)

Item	Description
Last Week Comp Ratio	The average compression ratio for the last seven days.

- Allows you to create, modify, and delete storage units.
- Displays four related tabs for a storage unit selected from the list: Storage Unit, Space Usage, Daily Written, and Data Movement.

  NOTE: The Data Movement tab is available only if an optional Cloud Tier license is installed.
- Takes you to **Replication** > **On-Demand** > **File Replication** when you click the **View DD Boost Replications** link.

  NOTE: A DD Replicator license is required for DD Boost to display tabs other than the File Replication tab.

Creating a storage unit

You must create at least one storage unit on the protection system, and a DD Boost user must be assigned to that storage unit. Use the Storage Units tab to create a storage unit.

About this task

Each storage unit is a top-level subdirectory of the `/data/col1` directory; there is no hierarchy among storage units.

Steps

1. Select **Protocols** > **DD Boost** > **Storage Units**.
2. Click **Create** (+).

   The Create Storage Unit dialog box is displayed.
3. Enter the storage unit name in the Name box.

   Each storage unit name must be unique. Storage unit names can be up to 50 characters. The following characters are acceptable:

   - upper- and lower-case alphabetical characters: A-Z, a-z
   - numbers: 0-9
   - embedded space

     NOTE: The storage-unit name must be enclosed in double quotes (") if the name has an embedded space.
   - comma (,)
   - period (.), as long as it does not precede the name
   - exclamation mark (!)
   - number sign (#)
   - dollar sign ($)
   - per cent sign (%)
   - plus sign (+)
   - at sign (@)
   - equal sign (=)
   - ampersand (&)
   - semi-colon (;)
   - parenthesis [(and)]
   - square brackets ([and])
   - curly brackets ({and})
   - caret (^)
   - tilde (~)
   - apostrophe (unslanted single quotation mark)
   - single slanted quotation mark (')
   - minus sign (-)
   - underscore (_)
4. To select an existing username that will have access to this storage unit, select the user name in the dropdown list.

If possible, select a username with management role privileges set to *none*.

5. To create and select a new username that will have access to this storage unit, select **Create a new Local User** and:
   a. Enter the new user name in the User box.

      The user must be configured in the backup application to connect to the protection system.

   b. Enter the password twice in the appropriate boxes.
6. To set storage space restrictions to prevent a storage unit from consuming excess space: enter either a soft or hard limit quota setting, or both a hard and soft limit. With a soft limit an alert is sent when the storage unit size exceeds the limit, but data can still be written to it. Data cannot be written to the storage unit when the hard limit is reached.

   NOTE: Quota limits are pre-compressed values. To set quota limits, select **Set to Specific Value** and enter the value. Select the unit of measurement: MiB, GiB, TiB, or PiB.

   NOTE: When setting both soft and hard limits, a quota's soft limit cannot exceed the quota's hard limit.

7. Click **Create**.
8. Repeat the above steps for each DD Boost-enabled system.

Viewing storage unit information

From the DD Boost Storage Units tab, you can select a storage unit and access the Storage Unit, Space Usage, Daily Written, and Data Movement tabs for the selected storage unit.

Storage Unit tab

The Storage Unit tab shows detailed information for a selected storage unit in its Summary and Quota panels. The Snapshot panel shows snapshot details, allows you to create new snapshots and schedules, and provides a link to the **Data Management** > **Snapshots** tab.

- The Summary panel shows summarized information for the selected storage unit.

Table 34. Summary panel

Summary item	Description
Total Files	The total number of file images on the storage unit. For compression details that you can download to a log file, click the Download Compression Details link. The generation can take up to several minutes. After it has completed, click Download.
Full Path	`/data/col1/filename`
Status	R: read; W: write; Q: quota defined
Pre-Comp Used	The amount of pre-compressed storage already used.

- The Quota panel shows quota information for the selected storage unit.

Table 35. Quota panel

Quota item	Description
Quota Enforcement	Enabled or disable. Clicking Quota takes you to the **Data Management** > **Quota** tab where you can configure quotas.
Pre-Comp Soft Limit	Current value of soft quota set for the storage unit.
Pre-Comp Hard Limit	Current value of hard quota set for the storage unit.
Quota Summary	Percentage of Hard Limit used.

To modify the pre-comp soft and hard limits shown in the tab:

1. Click the **Quota** link in the Quota panel.
2. In the Configure Quota dialog box, enter values for hard and soft quotas and select the unit of measurement: MiB, GiB, TiB, or PiB. Click **OK**.

- Snapshots

  The Snapshots panel shows information about the storage unit's snapshots.

  Table 36. Snapshots panel

Item	Description
Total Snapshots	The total number of snapshots created for this MTree. A total of 750 snapshots can be created for each MTree.
Expired	The number of snapshots in this MTree that have been marked for deletion, but have not been removed with the clean operation as yet.
Unexpired	The number of snapshots in this MTree that are marked for keeping.
Oldest Snapshot	The date of the oldest snapshot for this MTree.
Newest Snapshot	The date of the newest snapshot for this MTree.
Next Scheduled	The date of the next scheduled snapshot.
Assigned Snapshot Schedules	The name of the snapshot schedule assigned to this MTree.

  Using the Snapshots panel, you can:

  - Assign a snapshot schedule to a selected storage unit: Click **Assign Schedules**. Select the schedule's checkbox; click **OK** and **Close**.
  - Create a new schedule: Click **Assign Snapshot Schedules** > **Create Snapshot Schedule**. Enter the new schedule's name.

    NOTE: The snapshot name can be composed only of letters, numbers, **_**, **-**, **%d** (numeric day of the month: 01-31), **%a** (abbreviated weekday name), **%m** (numeric month of the year: 01-12), **%b** (abbreviated month name), **%y** (year, two digits), **%Y** (year, four digits), **%H** (hour: 00-23), and **%M** (minute: 00-59), following the pattern shown in the dialog box. Enter the new pattern and click **Validate Pattern & Update Sample**. Click **Next**.

    - Select when the schedule is to be executed: weekly, every day (or selected days), monthly on specific days that you select by clicking that date in the calendar, or on the last day of the month. Click **Next**.
    - Enter the times of the day when the schedule is to be executed: Either select At Specific Times or In Intervals. If you select a specific time, select the time from the list. Click Add (**+**) to add a time (24-hour format). For intervals, select In Intervals and set the start and end times and how often (Every), such as every eight hours. Click **Next**.
    - Enter the retention period for the snapshots in days, months, or years. Click **Next**.
    - Review the Summary of your configuration. Click **Back** to edit any of the values. Click **Finish** to create the schedule.
  - Click the Snapshots link to go to the **Data Management** > **Snapshots** tab.

Space Usage tab

The Space Usage tab graph displays a visual representation of data usage for the storage unit over time.

- Click a point on a graph line to display a box with data at that point.
- Click **Print** (at the bottom on the graph) to open the standard Print dialog box.
- Click **Show in new window** to display the graph in a new browser window.

There are two types of graph data displayed: Logical Space Used (Pre-Compression) and Physical Capacity Used (Post-Compression).

Daily Written tab

The Daily Written view contains a graph that displays a visual representation of data that is written daily to the system over a period of time, selectable from 7 to 120 days. The data amounts are shown over time for pre- and post-compression amounts.

Data Movement tab

A graph in the same format as the Daily Written graph that shows the amount of disk space moved to Cloud Tier storage (if the Cloud Tier license is enabled).

Modifying a storage unit

Use the Modify Storage Unit dialog to rename a storage unit, select a different existing user, create and select a new user, and edit quota settings.

About this task

Steps

1. Select **Protocols** > **DD Boost** > **Storage Units**.
2. In the Storage Unit list, select the storage unit to modify.
3. Click the pencil icon.

   The Modify Storage Unit dialog appears.
4. To rename the storage unit, edit the text in the **Name** field.
5. To select a different existing user, select the user name in the drop-down list.

   If possible, select a username with management role privileges set to *none*.
6. To create and select a new user, select **Create a new Local User** and do the following:
   a. Enter the new user name in the User box.

      The user must be configured in the backup application to connect to the protection system.

   b. Enter the password twice in the appropriate boxes.
7. Edit the Quota Settings as needed.

   To set storage space restrictions to prevent a storage unit from consuming excess space: enter either a soft or hard limit quota setting, or both a hard and soft limit. With a soft limit an alert is sent when the storage unit size exceeds the limit, but data can still be written to it. Data cannot be written to the storage unit when the hard limit is reached.

   NOTE: Quota limits are pre-compressed values. To set quota limits, select **Set to Specific Value** and enter the value. Select the unit of measurement: MiB, GiB, TiB, or PiB.

   NOTE: When setting both soft and hard limits, a quota's soft limit cannot exceed the quota's hard limit.
8. Click **Modify**.

Renaming a storage unit

Use the Modify Storage Unit dialog to rename a storage unit.

About this task

Renaming a storage unit changes the name of the storage unit while retaining its:

- Username ownership
- Stream limit configuration
- Capacity quota configuration and physical reported size
- AIR association on the local protection system

Steps

1. Go to **Protocols** > **DD Boost** > **Storage Units**.
2. In the Storage Unit list, select the storage unit to rename.

3. Click the pencil icon.

   The Modify Storage Unit dialog appears.

4. Edit the text in the **Name** field.
5. Click **Modify**.

Deleting a storage unit

Use the Storage Units tab to delete a storage unit from your protection system. Deleting a storage unit removes the storage unit, as well as any images contained in the storage unit, from your system.

Steps

1. Select **Protocols** > **DD Boost** > **Storage Units**.
2. Select the storage unit to be deleted from the list.
3. Click **Delete** (X).
4. Click **OK**.

Results

The storage unit is removed from your system. You must also manually remove the corresponding backup application catalog entries.

Undeleting a storage unit

Use the Storage Units tab to undelete a storage unit.

About this task

Undeleting a storage unit recovers a previously deleted storage unit, including its:

- Username ownership
- Stream limit configuration
- Capacity quota configuration and physical reported size
- AIR association on the local protection system

NOTE: Deleted storage units are available until the next `filesys clean` command is run.

Steps

1. Select **Protocols** > **DD Boost** > **Storage Units** > **More Tasks** > **Undelete Storage Unit...**.
2. In the Undelete Storage Units dialog box, select the storage unit(s) that you want to undelete.
3. Click **OK**.

Selecting DD Boost options

Use the Set DD Boost Options dialog to specify settings for distributed segment processing, virtual synthetics, low bandwidth optimization for file replication, file replication encryption, and file replication network preference (IPv4 or IPv6).

Steps

1. To display the DD Boost option settings, select **Protocols** > **DD Boost** > **Settings** > **Advanced Options**.
2. To change the settings, select **More Tasks** > **Set Options**.

   The Set DD Boost Options dialog appears.

3. Select any option to be enabled.
4. Deselect any option to be disabled.

   To deselect a File Replication Network Preference option, select the other option.

5. Set the DD Boost security options.
   a. Select the **Authentication Mode**:
      - None
      - Two-way
      - Two-way Password
   b. Select the **Encryption Strength**:

      NOTE: The default encryption strength is Medium.

      - None
      - Medium
      - High

   The protection system compares the global authentication mode and encryption strength against the per-client authentication mode and encryption strength to calculate the effective authentication mode and authentication encryption strength. The system does not use the highest authentication mode from one entry, and the highest encryption settings from a different entry. The effective authentication mode and encryption strength come from the single entry that provides the highest authentication mode.
6. Set the replication retry settings.
   a. Specify the number of times to retry a replication job.
   b. Specify the interval, in seconds, between replication retry attempts.
7. Click **OK**.

   NOTE: You can also manage distributed segment processing via the `ddboost option` commands, which are described in detail in the *DDOS Command Reference Guide*.

Distributed segment processing

Distributed segment processing increases backup throughput in almost all cases by eliminating duplicate data transmission between the media server and the protection system.

You can manage distributed segment processing via the `ddboost option` commands, which are described in detail in the *DDOS Command Reference Guide*.

Virtual synthetics

A virtual synthetic full backup is the combination of the last full (synthetic or full) backup and all subsequent incremental backups. Virtual synthetics are enabled by default.

Low-bandwidth optimization

If you use file replication over a low-bandwidth network (WAN), you can increase replication speed by using low bandwidth optimization. This feature provides additional compression during data transfer. Low bandwidth compression is available to protection systems with an installed Replication license.

Low-bandwidth optimization, which is disabled by default, is designed for use on networks with less than 6 Mbps aggregate bandwidth. Do not use this option if maximum file system write performance is required.

NOTE: You can also manage low bandwidth optimization via the `ddboost file-replication` commands, which are described in detail in the *DDOS Command Reference Guide*.

File replication encryption

You can encrypt the data replication stream by enabling the DD Boost file replication encryption option.

NOTE: If DD Boost file replication encryption is used on systems without the Data at Rest option, it must be set to on for both the source and destination systems.

Managed file replication TCP port setting

For DD Boost managed file replication, use the same global listen port on both the source and target protection systems. To set the listen port, use the `replication option` command as described in the *DDOS Command Reference Guide*.

File replication network preference

Use this option to set the preferred network type for DD Boost file replication to either IPv4 or IPv6.

Replication retry

By default, the system will retry a replication job 20 times before reporting it as failed, with a 30 second interval between retries. Configure these values to match the needs of the environment.

The number of retry attempts can be between 0 and 1000.

The retry interval can be between 1 and 3600 seconds.

Managing certificates for DD Boost

A host certificate allows DD Boost client programs to verify the identity of the system when establishing a connection. CA certificates identify certificate authorities that should be trusted by the system. The topics in this section describe how to manage host and CA certificates for DD Boost.

Adding a host certificate for DD Boost

Add a host certificate to your system. DDOS supports one host certificate for DD Boost.

Steps

1. If you have not yet requested a host certificate, request one from a trusted CA.
2. When you have received a host certificate, copy or move it to the computer from which you run DD Service Manager.
3. Start DD System Manager on the system to which you want to add a host certificate.

   NOTE: DD System Manager supports certificate management only on the management system (which is the system running DD System Manager).

4. Select **Protocols** > **DD Boost** > **More Tasks** > **Manage Certificates...**.

   NOTE: If you try to remotely manage certificates on a managed system, DD System Manager displays an information message at the top of the certificate management dialog. To manage certificates for a system, you must start DD System Manager on that system.

5. In the Host Certificate area, click **Add**.
6. To add a host certificate enclosed in a .p12 file, do the following:
   a. Select **I want to upload the certificate as a .p12 file**.
   b. Type the password in the **Password** box.
   c. Click **Browse** and select the host certificate file to upload to the system.
   d. Click **Add.**
7. To add a host certificate enclosed in a .pem file, do the following:
   a. Select **I want to upload the public key as a .pem file and use a generated private key**.
   b. Click **Browse** and select the host certificate file to upload to the system.
   c. Click **Add.**

Adding CA certificates for DD Boost

Add a certificate for a trusted CA to your system. DDOS supports multiple certificates for trusted CAs.

Steps

1. Obtain a certificate for the trusted CA.
2. Copy or move the trusted CA certificate to the computer from which you run DD Service Manager.
3. Start DD System Manager on the system to which you want to add the CA certificate.

   NOTE: DD System Manager supports certificate management only on the management system (which is the system running DD System Manager).

4. Select **Protocols** > **DD Boost** > **More Tasks** > **Manage Certificates...**.

   NOTE: If you try to remotely manage certificates on a managed system, DD System Manager displays an information message at the top of the certificate management dialog. To manage certificates for a system, you must start DD System Manager on that system.

5. In the CA Certificates area, click **Add**.

   The Add CA Certificate for DD Boost dialog appears.

6. To add a CA certificate enclosed in a .pem file, do the following:
   a. Select **I want to upload the certificate as a .pem file**.
   b. Click **Browse**, select the host certificate file to upload to the system, and click **Open**.
   c. Click **Add.**
7. To add a CA certificate using copy and paste, do the following:
   a. Copy the certificate text to the clipboard using the controls in your operating system.
   b. Select **I want to copy and paste the certificate text**.
   c. Paste the certificate text in the box below the copy and paste selection.
   d. Click **Add.**

Managing DD Boost client access and encryption

Use the DD Boost Settings tab to configure which specific clients, or set of clients, can establish a DD Boost connection with the protection System and whether or not the client will use encryption. By default, the system is configured to allow all clients to have access, with medium encryption.

NOTE: Enabling in-flight encryption will impact system performance.

NOTE: DD Boost offers global authentication and encryption options to defend your system against man-in-the-middle (MITM) attacks. You specify authentication and encryption settings using the GUI, or CLI commands on the protection system. For details, see the *DD Boost for OpenStorage 3.4 Administration Guide*, and Adding a DD Boost client or the *DDOS Command Reference Guide*.

Adding a DD Boost client

Create an allowed DD Boost client and specify whether the client will use encryption.

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.
2. In the Allowed Clients section, click **Create (+).**

   The Add Allowed Client dialog appears.

3. Enter the hostname of the client.

   This can be a fully-qualified domain name (e.g. host1.example.com) or a hostname with a wildcard (e.g. *.example.com).

4. Select the Encryption Strength.

The options are None (no encryption), Medium (AES128-SHA1), or High (AES256-SHA1).

NOTE: The default encryption strength is Medium.

5. Select the Authentication Mode.

   The options are One Way, Two Way, Two Way Password, or Anonymous.

6. Click **OK**.

Modifying a DD Boost client

Change the name, encryption strength, and authentication mode of an allowed DD Boost client.

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.
2. In the Allowed Clients list, select the client to modify.
3. Click the **Edit** button, which displays a pencil icon.

   The Modify Allowed Client dialog appears.

4. To change the name of a client, edit the Client text.
5. To change the Encryption Strength, select the option.

   The options are None (no encryption), Medium (AES128-SHA1), or High (AES256-SHA1).

6. To change the Authentication Mode, select the option.

   The options are One Way, Two Way, or Anonymous.

7. Click **OK**.

Removing a DD Boost client

Delete an allowed DD Boost client.

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.
2. Select the client from the list.
3. Click **Delete (X)**.

   The Delete Allowed Clients dialog appears.

4. Confirm and select the client name. Click **OK**.

About interface groups

This feature lets you combine multiple Ethernet links into a group and register only one interface on the protection system with the backup application. The DD Boost Library negotiates with the system to obtain the best interface to send data. Load balancing provides higher physical throughput to the system.

Configuring an interface group creates a private network within the system, comprised of the IP addresses designated as a group. Clients are assigned to a single group, and the group interface uses load balancing to improve data transfer performance and increase reliability.

For example, in the Veritas NetBackup environment, media server clients use a single public network IP address to access the system. All communication with the system is initiated via this administered IP connection, which is configured on the NetBackup server.

If an interface group is configured, when the system receives data from the media server clients, the data transfer is load-balanced and distributed on all the interfaces in the group, providing higher input/output throughput, especially for customers who use multiple 1 GigE connections.

The data transfer is load-balanced based on the number of connections outstanding on the interfaces. Only connections for backup and restore jobs are load-balanced. Check the Active Connections for more information on the number of outstanding connections on the interfaces in a group.

Should an interface in the group fail, all the in-flight jobs to that interface are automatically resumed on healthy operational links (unbeknownst to the backup applications). Any jobs that are started subsequent to the failure are also routed to a healthy interface in the group. If the group is disabled or an attempt to recover on an alternate interface fails, the administered IP is used for recovery. Failure in one group will not utilize interfaces from another group.

Consider the following information when managing interface groups.

- The IP address must be configured on the system, and its interface enabled. To check the interface configuration, select **Hardware** > **Ethernet** > **Interfaces** page, and check for free ports. See the `net` chapter of the *DDOS Command Reference Guide* for information about configuring an IP address for an interface.
- You can use the `ifgroup` commands to manage interface groups; these commands are described in detail in the *DDOS Command Reference Guide*.
- Interface groups provide full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 ifgroup interfaces only. A client connected with IPv4 sees IPv4 ifgroup interfaces only. Individual ifgroups include all IPv4 addresses or all IPv6 addresses. For details, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide*.
- Configured interfaces are listed in Active Connections, on the lower portion of the Activities page.

NOTE: See Using DD Boost on HA systems for important information about using interface groups with HA systems.

The topics that follow describe how to manage interface groups.

Related concepts

Active Connections

Interfaces

IFGROUP supports physical and bonded interfaces.

An IFGROUP interface is a member of a single IFGROUP *<group-name>* and may consist of:

- Physical interface such as `eth0a`
- Bonded interface, created for link failover or link aggregation, such as `veth1`
- Bonded alias interface such as `eth0a:2` or `veth1:2`
- Bonded VLAN interface such as `eth0a.1` or `veth1.1`
- Within an IFGROUP *<group-name>*, all interfaces must be on unique interfaces (Ethernet, bonded Ethernet) to ensure failover in the event of network error.

IFGROUP provides full support for static IPv6 addresses, providing the same capabilities for IPv6 as for IPv4. Concurrent IPv4 and IPv6 client connections are allowed. A client connected with IPv6 sees IPv6 IFGROUP interfaces only. A client connected with IPv4 sees IPv4 IFGROUP interfaces only. Individual IFGROUPs include all IPv4 addresses or all IPv6 addresses.

For more information, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide*.

Interface enforcement

IFGROUP lets you enforce private network connectivity, ensuring that a failed job does not reconnect on the public network after network errors.

When interface enforcement is enabled, a failed job can only retry on an alternative private network IP address. Interface enforcement is only available for clients that use IFGROUP interfaces.

Interface enforcement is off (FALSE) by default. To enable interface enforcement, you must add the following setting to the system registry:

`system.ENFORCE_IFGROUP_RW=TRUE`

After you've made this entry in the registry, you must do a `filesys restart` for the setting to take effect.

For more information, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide*.

Clients

IFGROUP supports various naming formats for clients. Client selection is based on a specified order of precedence.

An IFGROUP client is a member of a single ifgroup *<group-name>* and may consist of:

- A fully qualified domain name (FQDN) such as ddboost.exampledomain.com
- A partial host, allowing search on the first *n* characters of the hostname. For example, when n=3, valid formats are `rtp_.*example.com` and `dur_.*example.com`. Five different values of *n* (1-5) are supported.
- Wild cards such as `*.exampledomain.com` or "*"
- A short name for the client, such as ddboost
- Client public IP range, such as 128.5.20.0/24

Prior to write or read processing, the client requests an IFGROUP IP address from the server. To select the client IFGROUP association, the client information is evaluated according to the following order of precedence.

1. IP address of the connected protection system. If there is already an active connection between the client and the system, and the connection exists on the interface in the IFGROUP, then the IFGROUP interfaces are made available for the client.
2. Connected client IP range. An IP mask check is done against the client source IP; if the client's source IP address matches the mask in the IFGROUP clients list, then the IFGROUP interfaces are made available for the client.
   - For IPv4, you can select five different range masks, based on network.
   - For IPv6, fixed masks /64, /112, and /128 are available.

   This host-range check is useful for separate VLANs with many clients where there isn't a unique partial hostname (domain).
3. Client Name: `abc-11.d1.com`
4. Client Domain Name: `*.d1.com`
5. All Clients: `*`

For more information, see the *DD Boost for Partner Integration Administration Guide*.

Creating interface groups

Use the IP Network tab to create interface groups and to add interfaces and clients to the groups.

About this task

Multiple interface groups improve the efficiency of DD Boost by allowing you to:

- Configure DD Boost to use specific interfaces configured into groups.
- Assign clients to one of those interface groups.
- Monitor which interfaces are active with DD Boost clients.

Create interface groups first, and then add clients (as new media servers become available) to an interface group.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Interface Groups section, click Add (**+**).
3. Enter the interface group name.
4. Select one or more interfaces. A maximum of 32 interfaces can be configured.

   NOTE: Depending upon aliasing configurations, some interfaces may not be selectable if they are sharing a physical interface with another interface in the same group. This is because each interface within the group must be on a different physical interface to ensure fail-over recovery.
5. Click **OK**.
6. In the Configured Clients section, click Add (**+**).
7. Enter a fully qualified client name or **`*.mydomain.com`**.

NOTE: The * client is initially available to the default group. The * client may only be a member of one ifgroup.

8. Select a previously configured interface group, and click **OK**.

Enabling and disabling interface groups

Use the IP Network tab to enable and disable interface groups.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Interface Groups section, select the interface group in the list.

   NOTE: If the interface group does not have both clients and interfaces assigned, you cannot enable the group.

3. Click **Edit** (pencil).
4. Click **Enabled** to enable the interface group; clear the checkbox to disable.
5. Click **OK**.

Modifying an interface group's name and interfaces

Use the IP Network tab to change an interface group's name and the interfaces associated with the group.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Interface Groups section, select the interface group in the list.
3. Click **Edit** (pencil).
4. Retype the name to modify the name.

   The group name must be one to 24 characters long and contain only letters, numbers, underscores, and dashes. It cannot be the same as any other group name and cannot be "default", "yes", "no", or "all."

5. Select or deselect client interfaces in the Interfaces list.

   NOTE: If you remove all interfaces from the group, it will be automatically disabled.

6. Click **OK**.

Deleting an interface group

Use the IP Network tab to delete an interface group. Deleting an interface group deletes all interfaces and clients associated with the group.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Interface Groups section, select the interface group in the list. The default group cannot be deleted.
3. Click Delete (**X**).
4. Confirm the deletion.

Adding a client to an interface group

Use the IP Network tab to add clients to interface groups.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Clients section, click Add (**+**).
3. Enter a name for the client.

   Client names must be unique and may consist of:
   - **FQDN**
   - ***.domain**
   - Client public IP range:
     - For IPv4, `xx.xx.xx.0/24` provides a 24-bit mask against the connecting IP. The /24 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.
     - For IPv6, `xxxx::0/112` provides a 112-bit mask against the connecting IP. The /112 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.

   Client names have a maximum length of 128 characters.
4. Select a previously configured interface group, and click **OK**.

Modifying a client's name or interface group

Use the IP Network tab to change a client's name or interface group.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Clients section, select the client.
3. Click **Edit** (pencil).
4. Type a new client name.

   Client names must be unique and may consist of:
   - **FQDN**
   - ***.domain**
   - Client public IP range:
     - For IPv4, `xx.xx.xx.0/24` provides a 24-bit mask against the connecting IP. The /24 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.
     - For IPv6, `xxxx::0/112` provides a 112-bit mask against the connecting IP. The /112 represents what bits are masked when the client's source IP address is evaluated for access to the IFGROUP.

   Client names have a maximum length of 128 characters.
5. Select a new interface group from the menu.

   NOTE: The old interface group is disabled if it has no clients.
6. Click **OK**.

Deleting a client from the interface group

Use the IP Network tab to delete a client from an interface group.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Clients section, select the client.
3. Click Delete (**X**).

   NOTE: If the interface group to which the client belongs has no other clients, the interface group is disabled.

4. Confirm the deletion.

Using interface groups for Managed File Replication (MFR)

Interface groups can be used to control the interfaces used for DD Boost MFR, to direct the replication connection over a specific network, and to use multiple network interfaces with high bandwidth and reliability for failover conditions. All protection system IP types are supported—IPv4 or IPv6, Alias IP/VLAN IP, and LACP/failover aggregation.

NOTE: Interface groups used for replication are different from the interface groups previously explained and are supported for DD Boost Managed File Replication (MFR) only. For detailed information about using interface groups for MFR, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide*.

Without the use of interface groups, configuration for replication requires several steps:

1. Adding an entry in the `/etc/hosts` file on the source system for the target system and hard coding one of the private LAN network interfaces as the destination IP address.
2. Adding a route on the source system to the target system specifying a physical or virtual port on the source system to the remote destination IP address.
3. Configuring LACP through the network on all switches between the systems for load balancing and failover.
4. Requiring different applications to use different names for the target system to avoid naming conflicts in the `/etc/hosts` file.

Using interface groups for replication simplifies this configuration through the use of the DDOS System Manager or DDOS CLI commands. Using interface groups to configure the replication path lets you:

- Redirect a hostname-resolved IP address away from the public network, using another private system IP address.
- Identify an interface group based on configured selection criteria, providing a single interface group where all the interfaces are reachable from the target system.
- Select a private network interface from a list of interfaces belonging to a group, ensuring that the interface is healthy.
- Provide load balancing across multiple system interfaces within the same private network.
- Provide a failover interface for recovery for the interfaces of the interface group.
- Provide host failover if configured on the source system.
- Use Network Address Translation (NAT)

The selection order for determining an interface group match for file replication is:

1. Local MTree (storage-unit) path and a specific remote system hostname
2. Local MTree (storage-unit) path with any remote system hostname
3. Any MTree (storage-unit) path with a specific system hostname

The same MTree can appear in multiple interface groups only if it has a different system hostname. The same system hostname can appear in multiple interface groups only if it has a different MTree path. The remote hostname is expected to be an FQDN, such as dd9900-1.example.com.

The interface group selection is performed locally on both the source system and the target system, independent of each other. For a WAN replication network, only the remote interface group needs to be configured since the source IP address corresponds to the gateway for the remote IP address.

Adding a replication path to an interface group

Use the IP Network tab to add replication paths to interface groups.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Replication Paths section, click Add (**+**).
3. Enter values for **MTree** and/or **Remote Host**.
4. Select a previously configured interface group, and click **OK**.

Modifying a replication path for an interface group

Use the IP Network tab to modify replication paths for interface groups.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Replication Paths section, select the replication path.
3. Click **Edit** (pencil).
4. Modify any or all values for **MTree**, **Remote Host**, or **Interface Group**.
5. Click **OK**.

Deleting a replication path for an interface group

Use the IP Network tab to delete replication paths for interface groups.

Steps

1. Select **Protocols** > **DD Boost** > **IP Network**.
2. In the Configured Replication Paths section, select the replication path.
3. Click Delete (**X**).
4. In the Delete Replication Path(s) dialog, click **OK**.

Destroying DD Boost

Use this option to permanently remove all of the data (images) contained in the storage units. When you disable or destroy DD Boost, the DD Boost FC service is also disabled. Only an administrative user can destroy DD Boost.

Steps

1. Manually remove (expire) the corresponding backup application catalog entries.

   NOTE: If multiple backup applications are using the same protection system, then remove all entries from each of those applications' catalogs.

2. Select **Protocols** > **DD Boost** > **More Tasks** > **Destroy DD Boost...**.
3. Enter your administrative credentials when prompted.
4. Click **OK**.

Configuring DD Boost-over-Fibre Channel

In earlier versions of DDOS, all communication between the DD Boost Library and any protection system was performed using IP networking. DDOS now offers Fibre Channel as an alternative transport mechanism for communication between the DD Boost Library and the system.

NOTE: Windows, Linux, HP-UX (64-bit Itanium architecture), AIX, and Solaris client environments are supported.

Enabling DD Boost users

Before you can configure the DD Boost-over-FC service on a protection system, you must add one or more DD Boost users and enable DD Boost.

Prerequisites

- Log in to DD System Manager. For instructions, see "Logging In and Out of DD System Manager."

  CLI equivalent

```
login as: sysadmin
Data Domain OS 5.7.x.x-12345
Using keyboard-interactive authentication.
Password:
```

- If you are using the CLI, ensure that the SCSI target daemon is enabled:

```
# scsitarget enable
Please wait ...
SCSI Target subsystem is enabled.
```

> **NOTE:** If you are using DD System Manager, the SCSI target daemon is automatically enabled when you enable the DD Boost-over-FC service (later in this procedure).

- Verify that the DD Boost license is installed. In DD System Manager, select **Protocols** > **DD Boost** > **Settings**. If the Status indicates that DD Boost is not licensed, click **Add License** and enter a valid license in the Add License Key dialog box.

  CLI equivalents

```
# elicense show
```

```
# elicense update license-file
```

Steps

1. Select **Protocols** > **DD Boost** > **Settings**.
2. In the Users with DD Boost Access section, specify one or more DD Boost user names.

   A DD Boost user is also a DDOS user. When specifying a DD Boost user name, you can select an existing DDOS user name, or you can create a new DDOS user name and make that name a DD Boost user. This release supports multiple DD Boost users. For detailed instructions, see "Specifying DD Boost User Names."

   CLI equivalents

```
# user add username [password password]
```

```
# ddboost set user-name exampleuser
```

3. Click **Enable** to enable DD Boost.

   CLI equivalent

```
# ddboost enable
Starting DDBOOST, please wait..............
DDBOOST is enabled.
```

Results

You are now ready to configure the DD Boost-over-FC service.

Configuring DD Boost

After you have added user(s) and enabled DD Boost, you need to enable the Fibre Channel option and specify the DD Boost Fibre Channel server name. Depending on your application, you may also need to create one or more storage units and install the DD Boost API/plug-in on media servers that will access the protection system.

Steps

1. Select **Protocols** > **DD Boost** > **Fibre Channel**.
2. Click **Enable** to enable Fibre Channel transport.

   CLI equivalent

```
# ddboost option set fc enabled
Please wait...
DD Boost option "FC" set to enabled.
```

3. To change the DD Boost Fibre Channel server name from the default (hostname), click **Edit**, enter a new server name, and click **OK**.

   CLI equivalent

```
# ddboost fc dfc-server-name set DFC-ddbeta2
DDBoost dfc-server-name is set to "DFC-ddbeta2" for DDBoost FC.
Configure clients to use "DFC-DFC-ddbeta2" for DDBoost FC.
```

4. Select **Protocols** > **DD Boost** > **Storage Units** to create a storage unit (if not already created by the application).

   You must create at least one storage unit on the system, and a DD Boost user must be assigned to that storage unit. For detailed instructions, see "Creating a Storage Unit."

   CLI equivalent

```
# ddboost storage-unit create storage_unit_name-su
```

5. Install the DD Boost API/plug-in (if necessary, based on the application).

   The DD Boost OpenStorage plug-in software must be installed on NetBackup media servers that need to access the system. This plug-in includes the required DD Boost Library that integrates with the system. For detailed installation and configuration instructions, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide*.

Results

You are now ready to verify connectivity and create access groups.

Verifying connectivity and creating access groups

Go to **Hardware** > **Fibre Channel** > **Resources** to manage initiators and endpoints for access points. Go to **Protocols** > **DD Boost** > **Fibre Channel** to create and manage DD Boost-over-FC access groups.

About this task

NOTE: Avoid making access group changes on a protection system during active backup or restore jobs. A change may cause an active job to fail. The impact of changes during active jobs depends on a combination of backup software and host configurations.

Steps

1. Select **Hardware** > **Fibre Channel** > **Resources** > **Initiators** to verify that initiators are present.

   It is recommended that you assign aliases to initiators to reduce confusion during the configuration process.

   CLI equivalent

```
# scsitarget initiator show list
Initiator      System Address                Group   Service
-----------    ----------------------   ----------   -------
initiator-1    21:00:00:24:ff:31:b7:16         n/a       n/a
initiator-2    21:00:00:24:ff:31:b8:32         n/a       n/a
initiator-3    25:00:00:21:88:00:73:ee         n/a       n/a
initiator-4    50:06:01:6d:3c:e0:68:14         n/a       n/a
initiator-5    50:06:01:6a:46:e0:55:9a         n/a       n/a
initiator-6    21:00:00:24:ff:31:b7:17         n/a       n/a
initiator-7    21:00:00:24:ff:31:b8:33         n/a       n/a
initiator-8    25:10:00:21:88:00:73:ee         n/a       n/a
initiator-9    50:06:01:6c:3c:e0:68:14         n/a       n/a
initiator-10   50:06:01:6b:46:e0:55:9a         n/a       n/a
```

```
tsm6_p23       21:00:00:24:ff:31:ce:f8   SetUp_Test       VTL
------------   -----------------------   ----------   -------
```

2. To assign an alias to an initiator, select one of the initiators and click the pencil (edit) icon. In the Name field of the Modify Initiator dialog, enter the alias and click **OK**.

   CLI equivalents

```
# scsitarget initiator rename initiator-1 initiator-renamed
Initiator 'initiator-1' successfully renamed.
```

```
# scsitarget initiator show list
Initiator           System Address                Group   Service
-----------------   -----------------------   ----------   -------
initiator-2         21:00:00:24:ff:31:b8:32         n/a       n/a
initiator-renamed   21:00:00:24:ff:31:b7:16         n/a       n/a
-----------------   -----------------------   ----------   -------
```

3. On the Resources tab, verify that endpoints are present and enabled.

   CLI equivalent

```
# scsitarget endpoint show list
-------------   --------------   ------------   -------   ------
endpoint-fc-0   5a               FibreChannel   Yes       Online
endpoint-fc-1   5b               FibreChannel   Yes       Online
-------------   --------------   ------------   -------   ------
```

4. Go to **Protocols** > **DD Boost** > **Fibre Channel**.
5. In the DD Boost Access Groups area, click the **+** icon to add an access group.
6. Enter a unique name for the access group. Duplicate names are not supported.

   CLI equivalent

```
# ddboost fc group create test-dfc-group
DDBoost FC Group "test-dfc-group" successfully created.
```

7. Select one or more initiators. Optionally, replace the initiator name by entering a new one. Click **Next.**

   CLI equivalent

```
#ddboost fc group add test-dfc-group initiator initiator-5
Initiator(s) "initiator-5" added to group "test-dfc-group".
```

   An initiator is a port on an HBA attached to a backup client that connects to the system for the purpose of reading and writing data using the Fibre Channel protocol. The WWPN is the unique World-Wide Port Name of the Fibre Channel port in the media server.

8. Specify the number of DD Boost devices to be used by the group. This number determines which devices the initiator can discover and, therefore, the number of I/O paths to the system. The default is one, the minimum is one, and the maximum is 64.

   CLI equivalent

```
# ddboost fc group modify Test device-set count 5
Added 3 devices.
```

   See the *DD Boost for OpenStorage Administration Guide* for the recommended value for different clients.

9. Indicate which endpoints to include in the group: all, none, or select from the list of endpoints. Click **Next**.

   CLI equivalents

```
# scsitarget group add Test device ddboost-dev8 primary-endpoint all
secondary-endpoint all
Device 'ddboost-dev8' successfully added to group.
```

```
# scsitarget group add Test device ddboost-dev8 primary-endpoint
endpoint-fc-1 secondary-endpoint fc-port-0
Device 'ddboost-dev8' is already in group 'Test'.
```

   When presenting LUNs via attached FC ports on HBAs, ports can be designated as primary, secondary or none. A primary port for a set of LUNs is the port that is currently advertizing those LUNs to a fabric. A secondary port is a port that will

broadcast a set of LUNs in the event of primary path failure (this requires manual intervention). A setting of none is used in the case where you do not wish to advertize selected LUNs. The presentation of LUNs is dependent upon the SAN topology.

10. Review the Summary and make any modifications. Click **Finish** to create the access group, which is displayed in the DD Boost Access Groups list.

    CLI equivalent

    ```
    # scsitarget group show detailed
    ```

    NOTE: To change settings for an existing access group, select it from the list and click the pencil icon (Modify).

Related concepts

Working with endpoints
Working with initiators
Working with a selected access group

Related tasks

Selecting endpoints for a device
Configuring the NDMP device TapeServer group

Deleting access groups

Use the Fibre Channel tab to delete access groups.

Steps

1. Select **Protocols** > **DD Boost** > **Fibre Channel**.
2. Select the group to be deleted from the DD Boost Access Groups list.

   NOTE: You cannot delete a group that has initiators assigned to it. Edit the group to remove the initiators first.

3. Click Delete (**X**).

Related concepts

Working with a selected access group

Related tasks

Selecting endpoints for a device
Configuring the NDMP device TapeServer group

Using DD Boost on HA systems

HA provides seamless failover of any application using DD Boost—that is, any backup or restore operation continues with no manual intervention required. All other DD Boost user scenarios are supported on HA systems as well, including managed file replication (MFR), distributed segment processing (DSP), filecopy, and dynamic interface groups (DIG).

Note these special considerations for using DD Boost on HA systems:

- On HA-enabled protection systems, failovers of the DD server occur in less than 10 minutes. However, recovery of DD Boost applications may take longer than this, because Boost application recovery cannot begin until the DD server failover is complete. In addition, Boost application recovery cannot start until the application invokes the Boost library.
- DD Boost on HA systems requires that the Boost applications be using Boost HA libraries; applications using non-HA Boost libraries do not see seamless failover.
- MFR will fail over seamlessly when both the source and destination systems are HA-enabled. MFR is also supported on partial HA configurations (that is, when either the source or destination system is enabled, but not both) when the failure occurs on the HA-enabled system. For more information, see the *DD Boost for OpenStorage Administration Guide* or the *DD Boost for Partner Integration Administration Guide*.

- Dynamic interface groups should not include IP addresses associated with the direct interconnection between the active and standby nodes.
- DD Boost clients must be configured to use floating IP addresses.

About the DD Boost tabs

Learn to use the DD Boost tabs in DD System Manager.

Settings

Use the Settings tab to enable or disable DD Boost, select clients and users, and specify advanced options.

The Settings tab shows the DD Boost status (Enabled or Disabled). Use the **Status** button to switch between **Enabled** or **Disabled**.

Under **Allowed Clients**, select the clients that are to have access to the system. Use the **Add**, **Modify**, and **Delete** buttons to manage the list of clients.

Under **Users with DD Boost Access**, select the users that are to have DD Boost access. Use the **Add**, **Change Password**, and **Remove** buttons to manage the list of users.

Expand **Advanced Options** to see which advanced options are enabled. Go to **More Tasks** > **Set Options** to reset these options.

Active Connections

Use the Active Connections tab to see information about clients, interfaces, and outbound files.

Table 37. Connected client information

Item	Description
Client	The name of the connected client.
Idle	Whether the client is idle (Yes) or not (No).
Plug-In Version	The DD Boost plug-in version installed, such as 7.0.0.1.
OS Version	The operating system version installed, such as Linux 3.0.101-108.57-default x86_64
Application Version	The backup application version installed, such as NetWorker 19.1.
Encrypted	Whether the connection is encrypted (Yes) or not (No).
DSP	Whether or not the connection is using Distributed Segment Processing (DSP) or not.
Transport	Type of transport being used, such as IPv4, IPv6 or DFC (Fibre Channel).

Table 38. Configured interface connection information

Item	Description
Interface	The IP address of the interface.
Interface Group	One of the following: ● The name of the interface group. ● None, if not a member of one.
Backup	The number of active backup connections.
Restore	The number of active restore connections.
Replication	The number of active replication connections.
Synthetic	The number of synthetic backups.

Table 38. Configured interface connection information (continued)

Item	Description
Total	The total number of connections for the interface.

Table 39. Outbound file replication information

Outbound files item	Description
File Name	The name of the outgoing image file.
Target Host	The name of the host receiving the file.
Logical Bytes to Transfer	The number of logical bytes to be transferred.
Logical Bytes Transferred	The number of logical bytes already transferred.
Low Bandwidth Optimization	The number of low-bandwidth bytes already transferred.

IP Network

The IP Network tab lists configured interface groups. Details include whether or not a group is enabled and any configured client interfaces. Administrators can use the Interface Group menu to view which clients are associated with an interface group.

Fibre Channel

The Fibre Channel tab lists configured DD Boost access groups. Use the Fibre Channel tab to create and delete access groups and to configure initiators, devices, and endpoints for DD Boost access groups.

Storage Units

Use the **Storage Units** tab to view, create, modify, and delete storage units.

Table 40. Storage Units tab

Item	Description
View DD Boost Replications	View DD Boost replication contexts.
Storage Unit	The name of the storage unit.
User	Username associated with the storage unit.
Quota Hard Limit	The hard quota set for the storage unit.
Last 24hr Pre-Comp	The amount of data written to the storage unit in the last 24 hours, before compression.
Last 24hr Post-Comp	The amount of data written to the storage unit in the last 24 hours, after compression.
Last 24hr Comp Ratio	Compression ratio of the data written to the storage unit in the last 24 hours.
Weekly Avg Post-Comp	Average amount of data written to the storage unit each week, after compression.
Last Week Post-Comp	Amount of data written to the storage unit in the last week, after compression.
Weekly Avg Comp Ratio	Average compression ratio of data written to the storage unit each week.
Last Week Comp Ratio	Compression ratio of the data written to the storage unit in the last week.

Select a storage unit to see detailed information about it. Detailed information is available on three tabs:

- Storage Unit tab

Table 41. Storage unit details: Storage Unit tab

Item	Description
Total Files	The total number of file images on the storage unit.
Full Path	The full path of the storage unit.
Status	The current status of the storage unit (combinations are supported). Status can be: ○ D—Deleted ○ RO—Read-only ○ RW—Read/write ○ RD—Replication destination ○ RLE—DD Retention lock enabled ○ RLD—DD Retention lock disabled
Pre-Comp Used	The amount of pre-compressed storage already used.
Used (Post-Comp)	The total size after compression of the files in the storage unit.
Compression	The compression ratio achieved on the files.
Schedules	The number of physical capacity mesaurement schedules assigned to the storage unit.
Submitted Measurements	The number of times the physical capacity of the storage unit has been measured.
Quota Enforcement	Click Quota to go to the Data Management Quota page, which lists hard and soft quota values/percentage used by MTrees.
Pre-Comp Soft Limit	Current value of soft quota set for the storage unit.
Pre-Comp Hard Limit	Current value of hard quota set for the storage unit.
Quota Summary	Percentage of Hard Limit used.
Total Snapshots	Total number of snapshots of the storage unit.
Expired	Number of expired snapshots of the storage unit.
Unexpired	Number of unexpired snapshots of the storage unit.
Oldest Snapshot	The oldest snapshot of the storage unit.
Newest Snapshot	The newest snapshot of the storage unit.
Next Scheduled	The next scheduled snapshot of the storage unit.
Assigned Snapshot Schedules	The snapshot schedules assigned to the storage unit.

- Space Usage tab: Displays a graph showing pre-compression bytes used, post-compression bytes used, and compression factor.
- Daily Written tab: Displays a graph showing pre-compression bytes written, post-compression bytes written, and total compression factor.

22

DD Virtual Tape Library

This chapter includes:

Topics:

DD Virtual Tape Library overview

DD Virtual Tape Library (DD VTL) is a disk-based backup system that emulates the use of physical tapes. It enables backup applications to connect to and manage DD system storage using functionality almost identical to a physical tape library.

Virtual tape drives are accessible to backup software in the same way as physical tape drives. After you create these drives in a DD VTL, they appear to the backup software as SCSI tape drives. The DD VTL, itself, appears to the backup software as a SCSI robotic device accessed through standard driver interfaces. However, the backup software (not the DD system that is configured as a DD VTL) manages the movement of the media changer and backup images.

The following terms have special meaning when used with DD VTL:

- *Library*: A library emulates a physical tape library with drives, changer, CAPs (cartridge access ports), and slots (cartridge slots).
- *Tape*: A tape is represented as a file. Tapes can be imported from the vault to a library. Tapes can be exported from a library to the vault. Tapes can be moved within a library across drives, slots, and CAPs.
- *Pool*: A pool is a collection of tapes that maps to a directory on the file system. Pools are used to replicate tapes to a destination. By default, pools are created as MTree pools unless you specify them as directory pools when they are created. You can convert directory-based pools to MTree-based pools to take advantage of the greater functionality of MTrees.
- *Vault*: The vault holds tapes not being used by any library. Tapes reside in either a library or the vault.

DD VTL has been tested with, and is supported by, specific backup software and hardware configurations. For more information, see E-Lab Interoperability Navigator at https://elabnavigator.dell.com/eln/elnhome.

DD VTL supports simultaneous use of the tape library and file system (NFS/CIFS/DD Boost) interfaces.

When DR (disaster recovery) is needed, pools and tapes can be replicated to a remote DD system using the DD Replicator.

To protect data on tapes from modification, tapes can be locked using DD Retention Lock Governance software.

NOTE: At present, 16 Gb/s is supported for fabric and point-to-point topologies. Other topologies will present issues.

The KB articles *Data Domain: VTL Best Practices Guide* and *Data Domain: Create a Virtual Tape Library via CLI* , available from the Online Support website, provide additional information.

Related tasks

Changing a tape's write or retention lock state
Creating an MTree or pool replication pair

Planning a DD VTL

The DD VTL (Virtual Tape Library) feature has very specific requirements, such as proper licensing, interface cards, user permissions, etc. These requirements are listed here, complete with details and recommendations.

- An appropriate DD VTL license.
  - DD VTL is a licensed feature, and you must use NDMP (Network Data Management Protocol) over IP (Internet Protocol) or DD VTL directly over FC (Fibre Channel).
  - An additional license is required for IBM i systems – the I/OS license.
  - Adding a DD VTL license through the DD System Manager automatically disables and enables the DD VTL feature.
- An installed FC interface card or DD VTL configured to use NDMP.
  - If the DD VTL communication between a backup server and a DD system is through an FC interface, the DD system must have an FC interface card installed. Notice that whenever an FC interface card is removed from (or changed within) a DD system, any DD VTL configuration associated with that card must be updated.
  - If the DD VTL communication between a backup server and a DD system is through NDMP, no FC interface card is required. However, you must configure the TapeServer access group. Also, when using NDMP, all initiator and port functionality does not apply.
  - The net filter must be configured to allow the NDMP client to send information to the DD system. Run the **net filter add operation allow clients *<client-IP-address>*** command to allow access for the NDMP client.
    - For added security, run the **net filter add operation allow clients *<client-IP-address>* interfaces *<DD-interface-IP-address>*** command.
    - Add the **seq-id 1** option to the command to enforce this rule before any other net filter rules.
- A backup software minimum record (block) size.
  - If possible, set backup software to use a minimum record (block) size of 64 KiB or larger. Larger sizes usually give faster performance and better data compression.
  - Depending on your backup application, if you change the size after the initial configuration, data written with the original size might become unreadable.
- Appropriate user access to the system.
  - For basic tape operations and monitoring, only a user login is required.
  - To enable and configure DD VTL services and perform other configuration tasks, a sysadmin login is required.

Related tasks

Configuring the NDMP device TapeServer group

DD VTL limits

Before setting up or using a DD VTL, review these limits on size, slots, etc.

- I/O Size – The maximum supported I/O size for any DD system using DD VTL is 1 MB.
- Libraries – DD VTL supports a maximum of 64 libraries per DD system (that is, 64 DD VTL instances on each DD system).
- Initiators – DD VTL supports a maximum of 1024 initiators or WWPNs (world-wide port names) per DD system.
- Tape Drives – Information about tape drives is presented in the next section.
- Data Streams – Information about data streams is listed in Best practices for data streams sent to DD systems .
- Slots – DD VTL supports a maximum of:
  - 32,000 slots per library
  - 64,000 slots per DD system

  The DD system automatically adds slots to keep the number of slots equal to, or greater than, the number of drives.

  NOTE: Some device drivers (for example, IBM AIX atape device drivers) limit library configurations to specific drive/slot limits, which may be less than what the DD system supports. Backup applications, and drives used by those applications, may be affected by this limitation.
- CAPs (cartridge access ports) – DD VTL supports a maximum of:
  - 100 CAPs per library

- 1000 CAPs per DD system

Related concepts

Best practices for data streams sent to DD systems

Number of drives supported by a DD VTL

The maximum number of drives supported by a DD VTL depends on the number of CPU cores and the amount of memory installed (both RAM and NVRAM, if applicable) on a DD system.

NOTE: There are no references to model numbers in this table because there are many combinations of CPU cores and memories for each model, and the number of supported drives depends *only* on the CPU cores and memories – not on the particular model, itself.

Table 42. Number of drives supported by a DD VTL

Number of CPU cores	RAM (in GB)	NVRAM (in GB)	Maximum number of supported drives
Fewer than 32	4 or less	NA	64
	More than 4, up to 38	NA	128
	More than 38, up to 128	NA	256
	More than 128	NA	540
32 to 39	Up to 128	Less than 4	270
	Up to 128	4 or more	540
	More than 128	NA	540
40 to 59	NA	NA	540
60 or more	NA	NA	1885

Related concepts

Working with access groups

Related tasks

Creating libraries
Creating drives

Tape barcodes

When you create a tape, you must assign a unique *barcode* (never duplicate barcodes as this can cause unpredictable behavior). Each barcode consists of eight characters: the first six are numbers or uppercase letters (0-9, A-Z), and the last two are the tape code for the supported tape type, as shown in the following table.

NOTE: Although a DD VTL barcode consists of eight characters, either six or eight characters may be transmitted to a backup application, depending on the changer type.

Table 43. Tape Codes by Tape Type

Tape Type	Default Capacity (unless noted)	Tape Code
LTO-1	100 GiB	L1
LTO-1	50 GiB (non-default)	LA[a]
LTO-1	30 GiB (non-default)	LB
LTO-1	10 GiB (non-default)	LC
LTO-2	200 GiB	L2

Table 43. Tape Codes by Tape Type (continued)

Tape Type	Default Capacity (unless noted)	Tape Code
LTO-3	400 GiB	L3
LTO-4	800 GiB	L4
LTO-5	1.5 TiB	L5
LTO-7 (default)	6 TiB	L7
LTO-8	12 TiB	L8

a. For TSM, use the L2 tape code if the LA code is ignored.

For multiple tape libraries, barcodes are automatically incremented, if the sixth character (just before the "L") is a number. If an overflow occurs (9 to 0), numbering moves one position to the left. If the next character to increment is a letter, incrementation stops. Here are a few sample barcodes and how each will be incremented:

- 000000L1 creates tapes of 100 GiB capacity and can accept a count of up to 100,000 tapes (from 000000 to 99999).
- AA0000LA creates tapes of 50 GiB capacity and can accept a count of up to 10,000 tapes (from 0000 to 9999).
- AAAA00LB creates tapes of 30GiB capacity and can accept a count of up to 100 tapes (from 00 to 99).
- AAAAAALC creates one tape of 10 GiB capacity. Only one tape can be created with this name.
- AAA350L1 creates tapes of 100 GiB capacity and can accept a count of up to 650 tapes (from 350 to 999).
- 000AAALA creates one tape of 50 GiB capacity. Only one tape can be created with this name.
- 5M7Q3KLB creates one tape of 30 GiB capacity. Only one tape can be created with this name.

LTO tape drive compatibility

You may have different generations of LTO (Linear Tape-Open) technology in your setup; the compatibility between these generations is presented in tabular form.

In this table:

- RW = read and write compatible
- R = read-only compatible
- — = not compatible

Table 44. LTO tape drive compatibility

tape format	LTO-8 drive	LTO-7 drive	LTO-5 drive	LTO-4 drive	LTO-3 drive	LTO-2 drive	LTO-1 drive
LTO-8 tape	RW	-	—	—	—	—	—
LTO-7 tape	R	RW	—	—	—	—	—
LTO-5 tape	—	R	RW	—	—	—	—
LTO-4 tape	—	—	RW	RW	—	—	—
LTO-3 tape	—	—	R	RW	RW	—	—
LTO-2 tape	—	—	—	R	RW	RW	—
LTO-1 tape	—	—	—	—	R	RW	RW

Setting up a DD VTL

To set up a simple DD VTL, use the Configuration Wizard, which is described in the *Getting Started* chapter.

Then, continue with the following topics to enable the DD VTL, create libraries, and create and import tapes.

NOTE: If the deployment environment includes an AS400 system as a DD VTL client, refer to Configuring DD VTL default options to configure the serial number prefix for VTL changers and drives before configuring the DD VTL relationship between the protection system and the AS400 client system.

Related tasks

Using the system configuration wizard

HA systems and DD VTL

HA systems are compatible with DD VTL; however, if a DD VTL job is in progress during a failover, the job will need to be restarted manually after the failover is complete.

The *DD Operating System Backup Compatibility Guide* provides additional details about the HBA, switch, firmware, and driver requirements for using DD VTL in an HA environment.

DD VTL tape out to cloud

DD VTL supports storing the VTL vault on Cloud Tier storage. To use this functionality, the protection system must be a supported Cloud Tier configuration, and have a Cloud Tier license in addition to the VTL license.

Configure and license the Cloud Tier storage before configuring DD VTL to use cloud storage for the vault. Cloud Tier provides additional information about the requirements for Cloud Tier, and how to configure Cloud Tier.

The FC and network interface requirements for VTL are the same for both cloud-based and local vault storage. DD VTL does not require special configuration to use cloud storage for the vault. When configuring the DD VTL, select the cloud storage as the vault location. However, when working with a cloud-based vault, there are some data management options that are unique to the cloud-based vault. Working with the cloud-based vault provides more information.

Managing a DD VTL

You can manage a DD VTL using the DD System Manager or the CLI. After you login, you can check the status of your DD VTL process, check your license information, and review and configure options.

Logging In

To use a graphical user interface (GUI) to manage your DD Virtual Tape Library (DD VTL), log in to the DD System Manager.

CLI Equivalent

You can also log in at the CLI:

```
login as: sysadmin
Data Domain OS
Using keyboard-interactive authentication.
Password:
```

Enabling SCSI Target Daemon (CLI only)

If you do log in from the CLI, you must enable the scsitarget daemon (the Fibre Channel service). This daemon is enabled during the DD VTL or DD Boost-FC enable selections in DD System Manager. In the CLI, these processes need to be enabled separately.

```
# scsitarget enable
Please wait ...
SCSI Target subsystem is enabled.
```

Accessing DD VTL

From the menu at the left of the DD System Manager, select **Protocols** > **VTL**.

Status

In the **Virtual Tape Libraries** > **VTL Service** area, you can see the status of your DD VTL process is displayed at the top, for example, **Enabled: Running**. The first part of the status will be **Enabled** (on) or **Disabled** (off). The second part will be one of the following process states.

Table 45. DD VTL process states

State	Description
Running	DD VTL process is enabled and active (shown in green).
Starting	DD VTL process is starting.
Stopping	DD VTL process is being shut down.
Stopped	DD VTL process is disabled (shown in red).
Timing out	DD VTL process crashed and is attempting an automatic restart.
Stuck	After several failed automatic restarts, the DD VTL process is unable to shut down normally, so an attempt is being made to kill it.

DD VTL License

The VTL License line tells you whether your DD VTL license has been applied. If it says Unlicensed, select **Add License**. Enter your license key in the Add License Key dialog. Select **Next** and **OK**.

NOTE: All license information should have been populated as part of the factory configuration process; however, if DD VTL was purchased later, the DD VTL license key may not have been available at that time.

CLI Equivalent

You can also verify that the DD VTL license has been installed at the CLI:

```
# elicense show
## License Key                    Feature
--    -------------------        -----------
1     DEFA-EFCD-FCDE-CDEF        Replication
2     EFCD-FCDE-CDEF-DEFA        VTL
--    -------------------        -----------
```

If the license is not present, each unit comes with documentation – a quick install card – which will show the licenses that have been purchased. Enter the following command to populate the license key.

```
# elicense update <license-file>
```

I/OS License (for IBM i users)

For customers of IBM i, the I/OS License line tells you whether your I/OS license has been applied. If it says Unlicensed, select **Add License**. You must enter a valid I/OS license in either of these formats: **xxxx-xxxx-xxxx-xxxx** or **xxxx-xxxx-xxxx-xxxx-xxxx**. Your I/OS license must be installed before creating a library and drives to be used on an IBM i system. Select **Next** and **OK**.

Related concepts

Working with access groups
Working with resources
Working with pools

Enabling DD VTL

Enabling DD VTL broadcasts the WWN of the protection system HBA to customer fabric and enables all libraries and library drives. If a forwarding plan is required in the form of change control processes, this process should be enabled to facilitate zoning.

Steps

1. Make sure that you have a DD VTL license and that the file system is enabled.
2. Select **Virtual Tape Libraries** > **VTL Service**.
3. To the right of the Status area, select **Enable**.
4. In the Enable Service dialog box, select **OK**.
5. After DD VTL has been enabled, note that Status will change to **Enabled: Running** in green. Also note that the configured DD VTL options are displayed in the Option Defaults area.

   CLI Equivalent

```
# vtl enable
Starting VTL, please wait ...
VTL is enabled.
```

Disabling DD VTL

Disabling DD VTL closes all libraries and shuts down the DD VTL process.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service**.
2. To the right of the Status area, select **Disable**.
3. In the Disable Service dialog, select **OK**.
4. After DD VTL has been disabled, notice that the Status has changed to **Disabled: Stopped** in red.

   CLI Equivalent

```
# vtl disable
```

DD VTL option defaults

The Option Default area of the VTL Service page displays the current settings for default DD VTL options (auto-eject, auto-offline, and barcode-length) that you can configure.

In the **Virtual Tape Libraries** > **VTL Service** area, the current default options for your DD VTL are displayed. Select **Configure** to change any of these values.

Table 46. Option Defaults

Item	Description
Property	Lists the configured options: ● auto-eject ● auto-offline ● barcode-length
Value	Provides the value for each configured option: ● auto-eject: default (disabled), enabled, or disabled ● auto-offline: default (disabled), enabled, or disabled ● barcode-length: default (8), 6, or 8

Configuring DD VTL default options

You can configure DD VTL default options when you add a license, create a library, or any time thereafter.

About this task

NOTE: DD VTLs are assigned global options, by default, and those options are updated whenever global options change, unless you change them manually using this method.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service**.
2. In the Option Defaults area, select **Configure**. In the Configure Default Options dialog box, change any of the default options, and then click **OK**.

Table 47. DD VTL default options

Option	Values	Notes
auto-eject	default (disabled), enable, or disable	Enabling auto-eject causes any tape put into a CAP (cartridge access port) to automatically move to the virtual vault, unless: • the tape came from the vault, in which case the tape stays in the CAP. • an `ALLOW_MEDIUM_REMOVAL` command with a 0 value (false) has been issued to the library to prevent the removal of the medium from the CAP to the outside world.
auto-offline	default (disabled), enable, or disable	Enabling auto-offline takes a drive offline automatically before a tape move operation is performed.
barcode-length	default (8), 6 or 8 [automatically set to 6 for L180, RESTORER-L180, and DDVTL changer models]	Although a DD VTL barcode consists of 8 characters, either 6 or 8 characters may be transmitted to a backup application, depending on the changer type.

NOTE: To disable all of these service options, select **Reset to Factory**, and the values will be immediately reset to factory defaults.

Next steps

If the DD VTL environment contains an AS400 as a DD VTL client, configure the DD VTL option for serial-number-prefix manually before adding the AS400 to the DD VTL environment. This is required to avoid duplicate serial numbers when there are multiple protection systems using DD VTL. The serial-number-prefix value must:

- Be a unique six digit value such that no other DD VTL on any system in the environment has the same prefix number
- Not end with a zero

Configure this value only once during the deployment of the system and the configuration of DD VTL. It will persist with any future DDOS upgrades on the system. Setting this value does not require a DD VTL service restart. Any DD VTL library created after setting this value will use the new prefix for the serial number.

`CLI equivalent`

```
# vtl option set serial-number-prefix value
# vtl option show serial-number-prefix
```

Working with libraries

A library emulates a physical tape library with drives, changer, CAPs (cartridge access ports), and slots (cartridge slots). Selecting **Virtual Tape Libraries** > **VTL Service** > **Libraries** displays detailed information for all configured libraries.

Table 48. Library information

Item	Description
Name	The name of a configured library.
Drives	The number of drives configured in the library.
Slots	The number of slots configured in the library.
CAPs	The number of CAPs (cartridge access ports) configured in the library.

From the More Tasks menu, you can create and delete libraries, as well as search for tapes.

Creating libraries

DD VTL supports a maximum of 128 libraries per system, that is, 128 concurrently active virtual tape library instances on each DD system.

Prerequisites

If the deployment environment includes an AS400 system as a DD VTL client, refer to Configuring DD VTL default options to configure the serial number prefix for VTL changers and drives before creating the DD VTL library and configuring the DD VTL relationship between the protection system and the AS400 client system.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries**.
2. Select **More Tasks** > **Library** > **Create**
3. In the Create Library dialog, enter the following information:

Table 49. Create Library dialog

Field	User input
Library Name	Enter a name of from 1 to 32 alphanumeric characters.
Number of Drives	**NOTE:** The maximum number of drives supported by a DD VTL depends on the number of CPU cores and the amount of memory installed (both RAM and NVRAM, if applicable) on a DD system. Enter the number of drives from 1 to 98. The number of drives to be created will correspond to the number of data streams that will write to a library.
Drive Model	Select the desired model from the drop-down list: • IBM-LTO-1 • IBM-LTO-2 • IBM-LTO-3 • IBM-LTO-4 • IBM-LTO-5 • IBM-LTO-7 (default) • IBM-LTO-8 • HP-LTO-3 • HP-LTO-4 Do not mix drive types, or media types, in the same library. This can cause unexpected results and/or errors in the backup operation.
Number of Slots	Enter the number of slots in the library. Here are some things to consider: • The number of slots must be equal to or greater than the number of drives.

Table 49. Create Library dialog (continued)

Field	User input
	• You can have up to 32,000 slots per individual library • You can have up to 64,000 slots per system. • Try to have enough slots so tapes remain in the DD VTL and never need to be exported to a vault – to avoid reconfiguring the DD VTL and to ease management overhead. • Consider any applications that are licensed by the number of slots. As an example, for a standard 100-GB cartridge you might configure 5000 slots. This would be enough to hold up tp 500 TB (assuming reasonably compressible data).
Number of CAPs	(Optional) Enter the number of cartridge access ports (CAPs). • You can have up to 100 CAPs per library. • You can have up to 1000 CAPs per system. Check your particular backup software application documentation on the Online Support Site for guidance.
Changer Model Name	Select the desired model from the drop-down list: • L180 (default) • RESTORER-L180 • TS3500 • TS3500-L53 • TS4500 • I2000 • I6000 • DDVTL Check your particular backup software application documentation on the Online Support Site for guidance. Also refer to the DD VTL support matrix to see the compatibility of emulated libraries to supported software.
Options	
auto-eject	default (disabled), enable, disable
auto-offline	default (disabled), enable, disable
barcode-length	default (8), 6, 8 [automatically set to 6 for L180, RESTORER-L180, and DD VTL changer models]

4. Select **OK**.

   After the Create Library status dialog shows **Completed**, select **OK**.

   The new library appears under the Libraries icon in the VTL Service tree, and the options you have configured appear as icons under the library. Selecting the library displays details about the library in the Information Panel.

   Note that access to VTLs and drives is managed with Access Groups.

CLI Equivalent

```
# vtl add NewVTL model L180 slots 50 caps 5
This adds the VTL library, NewVTL. Use 'vtl show config NewVTL' to view it.

# vtl drive add NewVTL count 4 model IBM-LTO-3
This adds 4 IBM-LTO-3 drives to the VTL library, NewVTL.
```

Deleting libraries

When a tape is in a drive within a library, and that library is deleted, the tape is moved to the vault. However, the tape's pool does not change.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries**.
2. Select **More Tasks** > **Library** > **Delete**.
3. In the Delete Libraries dialog, select or confirm the checkbox of the items to delete:
   - The name of each library, or
   - Library Names, to delete all libraries
4. Select **Next**.
5. Verify the libraries to delete, and select **Submit** in the confirmation dialogs.
6. After the Delete Libraries Status dialog shows `Completed`, select **Close**. The selected libraries are deleted from the DD VTL.

CLI Equivalent

```
# vtl del OldVTL
```

Searching for tapes

You can use a variety of criteria – location, pool, and/or barcode – to search for a tape.

Steps

1. Select **Virtual Tape Libraries** or **Pools**.
2. Select the area to search (library, vault, pool).
3. Select **More Tasks** > **Tapes** > **Search**.
4. In the Search Tapes dialog, enter information about the tape(s) you want to find.

Table 50. Search Tapes dialog

Field	User input
Location	Specify a location, or leave the default (All).
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.

5. Select **Search**.

Working with a selected library

Selecting **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* displays detailed information for a selected library.

Table 51. Devices

Item	Description
Device	The elements in the library, such a drives, slots, and CAPs (cartridge access ports).
Loaded	The number of devices with media loaded.

Table 51. Devices (continued)

Item	Description
Empty	The number of devices with no media loaded.
Total	The total number of loaded and empty devices.

Table 52. Options

Property	Value
auto-eject	enabled or disabled
auto-offline	enabled or disabled
barcode-length	6 or 8

Table 53. Tapes

Item	Description
Pool	The name of the pool where the tapes are located.
Tape Count	The number of tapes in that pool.
Capacity	The total configured data capacity of the tapes in that pool, in GiB (Gibibytes, the base-2 equivalent of GB, Gigabytes).
Used	The amount of space used on the virtual tapes in that pool.
Average Compression	The average amount of compression achieved on the data on the tapes in that pool.

From the More Tasks menu, you can delete, rename, or set options for a library; create, delete, import, export, or move tapes; and add or delete slots and CAPs.

Related tasks

Deleting libraries

Creating tapes

You can create tapes in either a library or a pool. If initiated from a pool, the system first creates the tapes, then imports them to the library.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library* or **Vault** or **Pools** > **Pools >** *pool*.
2. Select **More Tasks** > **Tapes** > **Create**.
3. In the Create Tapes dialog, enter the following information about the tape:

Table 54. Create Tapes dialog

Field	User input
Library (if initiated from a library)	If a drop-down menu is enabled, select the library or leave the default selection.
Pool Name	Select the name of the pool in which the tape will reside, from the drop-down list. If no pools have been created, use the Default pool.
Number of Tapes	For a library, select from 1 to 20. For a pool, select from 1 to 100,000, or leave the default (20). [Although the number of supported tapes is unlimited, you can create no more than 100,000 tapes at a time.]
Starting Barcode	Enter the initial barcode number (using the format A99000LA).

Table 54. Create Tapes dialog (continued)

Field	User input
Tape Capacity	(optional) Specify the number of GiBs from 1 to 15000 for each tape (this setting overrides the barcode capacity setting). For efficient use of disk space, use 100 GiB or fewer.

4. Select **OK** and **Close**.

CLI Equivalent

```
# vtl tape add A00000L1 capacity 100 count 5 pool VTL_Pool
... added 5 tape(s)...
```

NOTE: You must auto-increment tape volume names in base10 format.

Related concepts

Tape barcodes

Deleting tapes

You can delete tapes from either a library or a pool. If initiated from a library, the system first exports the tapes, then deletes them. The tapes must be in the vault, not in a library. On a Replication destination DD system, deleting a tape is not permitted.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library* or **Vault** or **Pools** > **Pools >** *pool*.
2. Select **More Tasks** > **Tapes** > **Delete**.
3. In the Delete Tapes dialog, enter search information about the tapes to delete, and select **Search**:

Table 55. Delete Tapes dialog

Field	User input
Location	If there is a drop-down list, select a library, or leave the default **Vault** selection.
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode, or leave the default (*) to search for a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page – possible values are 15, 30, and 45.
Select all pages	Select the **Select All Pages** checkbox to select all tapes returned by the search query.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

4. Select the checkbox of the tape that should be deleted or the checkbox on the heading column to delete all tapes, and select **Next**.
5. Select **Submit** in the confirmation window, and select **Close**.

NOTE: After a tape is removed, the physical disk space used for the tape is not reclaimed until after a file system cleaning operation.

CLI Equivalent

```
# vtl tape del barcode [count count] [pool pool]
```

For example:

```
# vtl tape del A00000L1
```

NOTE: You can act on ranges; however, if there is a missing tape in the range, the action will stop.

Importing tapes

Importing a tape means that an existing tape will be moved from the vault to a library slot, drive, or cartridge access port (CAP).

About this task

The number of tapes you can import at one time is limited by the number of empty slots in the library, that is, you cannot import more tapes than the number of currently empty slots.

To view the available slots for a library, select the library from the stack menu. The information panel for the library shows the count in the Empty column.

- If a tape is in a drive, and the tape origin is known to be a slot, a slot is reserved.
- If a tape is in a drive, and the tape origin is unknown (slot or CAP), a slot is reserved.
- If a tape is in a drive, and the tape origin is known to be a CAP, a slot is not reserved. (The tape returns to the CAP when removed from the drive.)
- To move a tape to a drive, see the section on moving tapes, which follows.

Steps

1. You can import tapes using either step a. or step b.
   a. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*. Then, select **More Tasks** > **Tapes** > **Import**. In the Import Tapes dialog, enter search information about the tapes to import, and select **Search**:

Table 56. Import Tapes dialog

Field	User input
Location	If there is a drop-down list, select the location of the tape, or leave the default of **Vault**.
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Select Destination > Device	Select the destination device where the tape will be imported. Possible values are Drive, CAP, and Slot.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

   Based on the previous conditions, a default set of tapes is searched to select the tapes to import. If pool, barcode, or count is changed, select Search to update the set of tapes available from which to choose.

   b. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* **> Changer** > **Drives >** *drive* **> Tapes**. Select tapes to import by selecting the checkbox next to:
      - An individual tape, or
      - The **Barcode** column to select all tapes on the current page, or
      - The **Select all pages** checkbox to select all tapes returned by the search query.

   Only tapes showing Vault in the Location can be imported.

Select **Import from Vault**. This button is disabled by default and enabled only if all of the selected tapes are from the Vault.

2. From the Import Tapes: library view, verify the summary information and the tape list, and select **OK**.
3. Select **Close** in the status window.

CLI Equivalent

```
# vtl tape show pool VTL_Pool
Processing tapes....
Barcode  Pool     Location State    Size    Used (%)            Comp ModTime
-------- -------- -------- ----- ------- --------------- ---- -------------------
A00000L3 VTL_Pool vault    RW    100 GiB 0.0 GiB (0.00%) 0x   2010/07/16 09:50:41
A00001L3 VTL_Pool vault    RW    100 GiB 0.0 GiB (0.00%) 0x   2010/07/16 09:50:41
A00002L3 VTL_Pool vault    RW    100 GiB 0.0 GiB (0.00%) 0x   2010/07/16 09:50:41
A00003L3 VTL_Pool vault    RW    100 GiB 0.0 GiB (0.00%) 0x   2010/07/16 09:50:41
A00004L3 VTL_Pool vault    RW    100 GiB 0.0 GiB (0.00%) 0x   2010/07/16 09:50:41
-------- -------- -------- ----- ------- --------------- ---- -------------------
VTL Tape Summary
----------------
Total number of tapes:     5
Total pools:               1
Total size of tapes:       500 GiB
Total space used by tapes: 0.0 GiB
Average Compression:       0.0x

# vtl import NewVTL barcode A00000L3 count 5 pool VTL_Pool
... imported 5 tape(s)...

# vtl tape show pool VTL_Pool
Processing tapes....

VTL Tape Summary
----------------
Total number of tapes:     5
Total pools:               1
Total size of tapes:       500 GiB
Total space used by tapes: 0.0 GiB
Average Compression:       0.0x
```

Exporting tapes

Exporting a tape removes that tape from a slot, drive, or cartridge-access port (CAP) and sends it to the vault.

Steps

1. You can export tapes using either step a. or step b.
   a. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*. Then, select **More Tasks** > **Tapes** > **Export**. In the Export Tapes dialog, enter search information about the tapes to export, and select **Search**:

Table 57. Export Tapes dialog

Field	User input
Location	If there is a drop-down list, select the name of the library where the tape is located, or leave the selected library.
Pool	Select the name of the pool in which to search for the tape. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.

Table 57. Export Tapes dialog (continued)

Field	User input
Select all pages	Select the **Select All Pages** checkbox to select all tapes returned by the search query.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

b. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* **>** **Changer** > **Drives >** *drive* **>** **Tapes**. Select tapes to export by selecting the checkbox next to:

- An individual tape, or
- The **Barcode** column to select all tapes on the current page, or
- The **Select all pages** checkbox to select all tapes returned by the search query.

Only tapes with a library name in the Location column can be exported.

Select **Export from Library**. This button is disabled by default and enabled only if all of the selected tapes have a library name in the Location column.

2. From the Export Tapes: library view, verify the summary information and the tape list, and select **OK**.
3. Select **Close** in the status window.

CLI Equivalent

```
# vtl export NewVTL cap address 1 count 4
... exported 4 tape(s)...
```

Moving tapes between devices within a library

Tapes can be moved between physical devices within a library to mimic backup software procedures for physical tape libraries (which move a tape in a library from a slot to a drive, a slot to a CAP, a CAP to a drive, and the reverse). In a physical tape library, backup software never moves a tape outside the library. Therefore, the destination library cannot change and is shown only for clarification.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** > ***library***.

   Note that when started from a library, the Tapes panel allows tapes to be moved only between devices.

2. Select **More Tasks** > **Tapes** > **Move**.

   Note that when started from a library, the Tapes panel allows tapes to be moved only between devices.

3. In the Move Tape dialog, enter search information about the tapes to move, and select **Search**:

Table 58. Move Tape dialog

Field	User input
Location	Location cannot be changed.
Pool	Select a pool.
Barcode	Specify a unique barcode. or leave the default (*) to return a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

4. From the search results list, select the tape or tapes to move.
5. Do one of the following:
   a. Select the device from the Device list (for example, a slot, drive, or CAP), and enter a starting address using sequential numbers for the second and subsequent tapes. For each tape to be moved, if the specified address is occupied, the next available address is used.
   b. Leave the address blank if the tape in a drive originally came from a slot and is to be returned to that slot; or if the tape is to be moved to the next available slot.
6. Select **Next**.
7. In the Move Tape dialog, verify the summary information and the tape listing, and select **Submit**.
8. Select **Close** in the status window.

Adding slots

You can add slots from a configured library to change the number of storage elements.

About this task

NOTE: Some backup applications do not automatically recognize that slots have been added to a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*.
2. Select **More Tasks** > **Slots** > **Add**.
3. In the Add Slots dialog, enter the Number of Slots to add. The total number of slots in a library, or in all libraries on a system, cannot exceed 32,000 for a library and 64,000 for a system.
4. Select **OK** and **Close** when the status shows `Completed`.

Related tasks

Configuring the NDMP device TapeServer group

Deleting slots

You can delete slots from a configured library to change the number of storage elements.

About this task

NOTE: Some backup applications do not automatically recognize that slots have been deleted from a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Steps

1. If the slot that you want to delete contains cartridges, move those cartridges to the vault. The system will delete only empty, uncommitted slots.
2. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*.
3. Select **More Tasks** > **Slots** > **Delete**.
4. In the Delete Slots dialog, enter the Number of Slots to delete.
5. Select **OK** and **Close** when the status shows `Completed`.

Adding CAPs

You can add CAPs (cartridge access ports) from a configured library to change the number of storage elements.

About this task

NOTE: CAPs are used by a limited number of backup applications. See your application documentation to ensure that CAPs are supported.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*.
2. Select **More Tasks** > **CAPs** > **Add**.
3. In the Add CAPs dialog, enter the Number of CAPs to add. You can add from 1 to 100 CAPs per library and from 1 to 1,000 CAPs per system.
4. Select **OK** and **Close** when the status shows `Completed`.

Related tasks

Configuring the NDMP device TapeServer group

Deleting CAPs

You can delete CAPs (cartridge access ports) from a configured library to change the number of storage elements.

About this task

NOTE: Some backup applications do not automatically recognize that CAPs have been deleted from a DD VTL. See your application documentation for information on how to configure the application to recognize this type of change.

Steps

1. If the CAP that you want to delete contains cartridges, move those cartridges to the vault, or this will be done automatically.
2. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library*.
3. Select **More Tasks** > **CAPs** > **Delete**.
4. In the Delete CAPs dialog, enter the Number of CAPs to delete. You can delete a maximum of 100 CAPs per library or 1000 CAPs per system.
5. Select **OK** and **Close** when the status shows `Completed`.

Viewing changer information

There can be only one changer per DD VTL. The changer model you select depends on your specific configuration.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** .
2. Select a specific library.
3. If not expanded, select the plus sign (**+**) on the left to open the library, and select a Changer element to display the Changer information panel, which provides the following information.

Table 59. Changer information panel

Item	Description
Vendor	The name of the vendor who manufactured the changer
Product	The model name
Revision	The revision level

Table 59. Changer information panel (continued)

Item	Description
Serial Number	The changer serial number

Working with drives

Selecting **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* > **Drives** displays detailed information for all drives for a selected library.

About this task

Table 60. Drives information panel

Column	Description
Drive	The list of drives by name, where name is "Drive #" and # is a number between 1 and n representing the address or location of the drive in the list of drives.
Vendor	The manufacturer or vendor of the drive, for example, IBM.
Product	The product name of the drive, for example, ULTRIUM-TD5.
Revision	The revision number of the drive product.
Serial Number	The serial number of the drive product.
Status	Whether the drive is Empty, Open, Locked, or Loaded. A tape must be present for the drive to be locked or loaded.
Tape	The barcode of the tape in the drive (if any).
Pool	The pool of the tape in the drive (if any).

Tape and library drivers – To work with drives, you must use the tape and library drivers supplied by your backup software vendor that support the IBM LTO-1, IBM LTO-2, IBM LTO-3, IBM LTO-4, IBM LTO-5, IBM LTO-7 (default), HP-LTO-3, or HP-LTO-4 drives and the StorageTek L180 (default), RESTORER-L180, IBM TS3500, I2000, I6000, or DDVTL libraries. For more information, see the *Application Compatibility Matrices and Integration Guides* for your vendors. When configuring drives, also keep in mind the limits on backup data streams, which are determined by the platform in use.

LTO capacities – Because the DD system treats LTO drives as virtual drives, you can set a maximum capacity to 15 TiB (15000 GiB) for each drive type. The default capacities for each LTO drive type are as follows:

- LTO-1 drive: 100 GiB
- LTO-2 drive: 200 GiB
- LTO-3 drive: 400 GiB
- LTO-4 drive: 800 GiB
- LTO-5 drive: 1.5 TiB (1500 GiB)
- LTO-7 drive: 6 TiB (6000 GiB)
- LTO-8 drive: 12 TiB (12000 GiB)

Migrating LTO-1 tapes – You can migrate tapes from existing LTO-1 type VTLs to VTLs that include other supported LTO-type tapes and drives. The migration options are different for each backup application, so follow the instructions in the LTO tape migration guide specific to your application. To find the appropriate guide, go to the Online Support Site, and in the search text box, type in **LTO Tape Migration for VTLs**.

Tape full: Early warning – You will receive a warning when the remaining tape space is almost completely full, that is, greater than 99.9, but less than 100 percent. The application can continue writing until the end of the tape to reach 100 percent capacity.

From the More Tasks menu, you can create or delete a drive.

Creating drives

See the *Number of drives supported by a DD VTL* section to determine the maximum number of drives supported for your particular DD VTL.

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library***> Changer** > **Drives**.
2. Select **More Tasks** > **Drives** > **Create**.
3. In the Create Drive dialog, enter the following information:

Table 61. Create Drive dialog

Field	User input
Location	Select a library name, or leave the name selected.
Number of Drives	See the table in the *Number of Drives Supported by a DD VTL* section, earlier in this chapter.
Model Name	Select the model from the drop-down list. If another drive already exists, this option is inactive, and the existing drive type must be used. You cannot mix drive types in the same library. • IBM-LTO-1 • IBM-LTO-2 • IBM-LTO-3 • IBM-LTO-4 • IBM-LTO-5 • IBM-LTO-7 (default) • IBM-LTO-8 • HP-LTO-3 • HP-LTO-4

4. Select **OK**, and when the status shows `Completed`, select **OK**.

   The added drive appears in the Drives list.

Related concepts

Number of drives supported by a DD VTL

Related tasks

Configuring the NDMP device TapeServer group

Deleting drives

A drive must be empty before it can be deleted.

Steps

1. If there is a tape in the drive that you want to delete, remove the tape.
2. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries >** *library* **> Changer** > **Drives**.
3. Select **More Tasks** > **Drives** > **Delete**.
4. In the Delete Drives dialog, select the checkboxes of the drives to delete, or select the **Drive** checkbox to delete all drives.
5. Select **Next**, and after verifying that the correct drive(s) has been selected for deletion, select **Submit**.
6. When the Delete Drive Status dialog shows `Completed`, select **Close**.

   The drive will have been removed from the Drives list.

Working with a selected drive

Selecting **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* **> Drives >** *drive* displays detailed information for a selected drive.

Table 62. Drive Tab

Column	Description
Drive	The list of drives by name, where name is "Drive #" and # is a number between 1 and n representing the address or location of the drive in the list of drives.
Vendor	The manufacturer or vendor of the drive, for example, IBM.
Product	The product name of the drive, for example, ULTRIUM-TD5.
Revision	The revision number of the drive product.
Serial Number	The serial number of the drive product.
Status	Whether the drive is Empty, Open, Locked, or Loaded. A tape must be present for the drive to be locked or loaded.
Tape	The barcode of the tape in the drive (if any).
Pool	The pool of the tape in the drive (if any).

Table 63. Statistics Tab

Column	Description
Endpoint	The specific name of the endpoint.
Ops/s	The operations per second.
Read KiB/s	The speed of reads in KiB per second.
Write KiB/s	The speed of writes in KiB per second.

From the More Tasks menu, you can delete the drive or perform a refresh.

Related tasks

Deleting drives

Working with tapes

A tape is represented as a file. Tapes can be imported from the vault to a library. Tapes can be exported from a library to the vault. Tapes can be moved within a library across drives, slots (cartridge slots), and CAPs (cartridge access ports).

About this task

When tapes are created, they are placed into the vault. After they have been added to the vault, they can be imported, exported, moved, searched, or removed.

Selecting **Virtual Tape Libraries** > **VTL Service** > **Libraries** > *library* > **Tapes** displays detailed information for all tapes for a selected library.

Table 64. Tape description

Item	Description
Barcode	The unique barcode for the tape.
Pool	The name of the pool that holds the tape. The Default pool holds all tapes unassigned to a user-created pool.

Table 64. Tape description (continued)

Item	Description
Location	The location of the tape - whether in a library (and which drive, CAP, or slot number) or in the virtual vault.
State	The state of the tape: ● RW – Read-writable ● RL – Retention-locked ● RO – Readable only ● WP – Write-protected ● RD – Replication destination
Capacity	The total capacity of the tape.
Used	The amount of space used on the tape.
Compression	The amount of compression performed on the data on a tape.
Last Modified	The date of the last change to the tape's information. Modification times used by the system for age-based policies might differ from the last modified time displayed in the tape information sections of the DD System Manager.
Locked Until	If a DD Retention Lock deadline has been set, the time set is shown. If no retention lock exists, this value is `Not specified`.

From the information panel, you can import a tape from the vault, export a tape to the library, set a tape's state, create a tape, or delete a tape.

From the More Tasks menu, you can move a tape.

Related tasks

Importing tapes
Exporting tapes
Creating tapes
Deleting tapes
Moving tapes between devices within a library

Changing a tape's write or retention lock state

Before changing a tape's write or retention lock state, the tape must have been created and imported. DD VTL tapes follow the standard DD Retention Lock policy. After the retention period for a tape has expired, it cannot be written to or changed (however, it can be deleted).

Steps

1. Select **Virtual Tape Libraries** > **VTL Service** > **Libraries** > ***library*** > **Tapes**.
2. Select the tape to modify from the list, and select **Set State** (above the list).
3. In the Set Tape State dialog, select **Read-Writeable**, **Write-Protected**, or **Retention-Lock**.
4. If the state is Retention-Lock, either
   - enter the tape's expiration date in a specified number of days, weeks, months, years, or
   - select the calendar icon, and select a date from the calendar. The Retention-Lock expires at noon on the selected date.
5. Select **Next**, and select **Submit** to change the state.

Working with the vault

The vault holds tapes not being used by any library. Tapes reside in either a library or the vault.

Selecting **Virtual Tape Libraries** > **VTL Service** > **Vault** displays detailed information for the Default pool and any other existing pools in the vault.

Systems with Cloud Tier and DD VTL provide the option of storing the vault on cloud storage. DD VTL does not support the option to store the vault from an MTree replication destination on cloud storage.

Table 65. Pool Summary

Item	Description
Pool Count	The number of VTL pools.
Tape Count	The number of tapes in the pools.
Size	The total amount of space in the pools.
Logical Used	The amount of space used in the pools.
Compression	The average amount of compression in the pools.

The **Protection Distribution** pane displays the following information.

NOTE: This table only appears if Cloud Tier is enabled on the protection system.

Table 66. Protection Distribution

Item	Description
Storage type	Vault or Cloud.
Cloud provider	For systems with tapes in Cloud Tier, there is a column for each cloud provider.
Logical Used	The amount of space used in the pools.
Pool Count	The number of VTL pools.
Tape Count	The number of tapes in the pools.

From the **More Tasks** menu, you can create, delete, and search for tapes in the vault.

Related tasks

Creating tapes
Deleting tapes
Searching for tapes

Working with the cloud-based vault

DD VTL supports several parameters that are unique to configurations where the vault is stored on Cloud Tier storage.

The following operations are available for working with cloud-based vault storage.

- Configure the data movement policy and cloud unit information for the specified VTL pool. Run the **vtl pool modify *<pool-name>* data-movement-policy {user-managed | age-threshold *<days>* | none} to-tier {cloud} cloud-unit *<cloud-unit-name>*** command.

  The available data movement policies are:

  - User-managed: The administrator can set this policy on a pool, to manually select tapes from the pool for migration to the cloud tier. The tapes migrate to the cloud tier on the first data movement operation after the tapes are selected.
  - Age-threshold: The administrator can set this policy on a pool, to allow the DD VTL to automatically select tapes from the pool for migration to the cloud tier based on the age of the tape. The tapes are selected for migration within six hours after they meet the age threshold, and are migrated on the first data movement operation after the tapes are selected.
- Select a specified tape for migration to the cloud tier. Run the **vtl tape select-for-move barcode *<barcode>* [count *<count>*] pool *<pool>* to-tier {cloud}** command.
- Deselect a specified tape for migration to the cloud tier. Run the **vtl tape deselect-for-move barcode *<barcode>* [count *<count>*] pool *<pool>* to-tier {cloud}** command.
- Recall a tape from the cloud tier. Run the **vtl tape recall start barcode *<barcode>* [count *<count>*] pool *<pool>*** command.

  After the recall, the tape resides in a local DD VTL vault and must be imported to the library for access.

NOTE: Run the `vtl tape show` command at any time to check the current location of a tape. The tape location updates within one hour of the tape moving to or from the cloud tier.

Prepare the VTL pool for data movement

Set the data movement policy on the VTL pool to manage migration of VTL data from the local vault to Cloud Tier.

About this task

Data movement for VTL occurs at the tape volume level. Individual tape volumes or collections of tape volumes can be moved to the cloud tier but only from the vault location. Tapes in other elements of a VTL cannot be moved.

NOTE: The default VTL pool and vault , `/data/col1/backup` directories or legacy library configurations cannot be used for Tape out to Cloud.

Steps

1. Select **Protocols** > **DD VTL**.
2. Expand the list of pools, and select a pool on which to enable migration to Cloud Tier.
3. In the **Cloud Data Movement** pane, click **Create** under **Cloud Data Movement Policy**.
4. In the **Policy** drop-down list, select a data movement policy:
   - **Age of tapes in days**
   - **Manual selection**
5. Set the data movement policy details.
   - For **Age of tapes in days**, select an age threshold after which tapes are migrated to Cloud Tier, and specify a destination cloud unit.
   - For **Manual selection**, specify a destination cloud unit.
6. Click **Create**.

   NOTE: After creating the data movement policy, the **Edit** and **Clear** buttons can be used to modify or delete the data movement policy.

CLI equivalent

Steps

1. Set the data movement policy to user-managed or age-threshold

   NOTE: VTL pool and cloud unit names are case sensitive and commands will fail if the case is not correct.

   - To set the data movement policy to user-managed, run the following command: `vtl pool modify cloud-vtl-pool data-movement-policy user-managed to-tier cloud cloud-unit ecs-unit1`

     ```
     **   Any tapes that are already selected will be migrated on the next data-movement run.
     VTL data-movement policy is set to "user-managed" for VTL pool "cloud-vtl-pool".
     ```

   - To set the data movement policy to age-threshold, run the following command:

     NOTE: The minimum is 14 days, and the maximum is 182,250 days.

     ```
     vtl pool modify cloud-vtl-pool data-movement-policy age-threshold 14 to-tier cloud cloud-unit ecs-unit1
     ```

     ```
     **   Any tapes that are already selected will be migrated on the next data-movement run.
     VTL data-movement policy "age-threshold" is set to 14 days for the VTL pool "cloud-vtl-pool".
     ```

2. Verify the data movement policy for the VTL pool.

   Run the following command: `vtl pool show all`

```
VTL Pools
Pool                    Status   Tapes   Size (GiB)   Used (GiB)   Comp   Cloud Unit
Cloud Policy
---------------------   ------   -----   ----------   ----------   ----   ----------
------------
cloud-vtl-pool          RW          50          250           41    45x   ecs-unit1
user-managed
Default                 RW           0            0            0     0x   -
none
---------------------   ------   -----   ----------   ----------   ----   ----------
------------
8080 tapes in 5 pools

 RO  : Read Only
 RD  : Replication Destination
 BCM : Backwards-Compatibility
```

3. Verify the policy for the VTL pool MTree is `app-managed`.

   Run the following command: `data-movement policy show all`

```
Mtree                       Target(Tier/Unit Name)   Policy        Value
-------------------------   ----------------------   -----------   -------
/data/col1/cloud-vtl-pool   Cloud/ecs-unit1          app-managed   enabled
-------------------------   ----------------------   -----------   -------
```

Remove tapes from the backup application inventory

Use the backup application verify the tape volumes that will move to the cloud are marked and inventoried according to the backup application requirements.

Select tape volumes for data movement

Manually select tapes for migration to Cloud Tier (immediately or at the next scheduled data migration), or manually remove tapes from the migration schedule.

Prerequisites

Verify the backup application is aware of status changes for volumes moved to cloud storage. Complete the necessary steps for the backup application to refresh its inventory to reflect the latest volume status.

If the tape is not in the vault, it cannot be migrated to Cloud Tier.

About this task

Steps

1. Select **Protocols** > **DD VTL**.
2. Expand the list of pools, and select the pool which is configured to migrate tapes to Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select tapes for migration to Cloud Tier.
5. Click **Select for Cloud Move** to migrate the tape at the next scheduled migration, or **Move to Cloud Now** to immediately migrate the tape.

   NOTE: If the data movement policy is based on tape ages, the **Select for Cloud Move** is not available, as the protection system automatically selects tapes for migration.

6. Click **Yes** at the confirmation dialog.

Unselect tape volumes for data movement

About this task

Tapes selected for migration to Cloud Tier can be removed from the migration schedule.

Steps

1. Select **Protocols** > **DD VTL**.
2. Expand the list of pools, and select the pool which is configured to migrate tapes to Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select tapes for migration to Cloud Tier.
5. Click **Unselect Cloud Move** to remove the tape from the migration schedule.
6. Click **Yes** at the confirmation dialog.

CLI equivalent

Steps

1. Identify the slot location of the tape volume to move.

   Run the following command: `vtl tape show cloud-vtl`

```
Processing tapes....
Barcode    Pool             Location            State    Size           Used (%)    Comp
 Modification Time
--------   --------------   -----------------   -----   -----   ----------------   ----
 -------------------
T00001L3   cloud-vtl-pool   cloud-vtl slot 1    RW      5 GiB   5.0 GiB (99.07%)   205x
 2017/05/05 10:43:43
T00002L3   cloud-vtl-pool   cloud-vtl slot 2    RW      5 GiB   5.0 GiB (99.07%)    36x
 2017/05/05 10:45:10
T00003L3   cloud-vtl-pool   cloud-vtl slot 3    RW      5 GiB   5.0 GiB (99.07%)    73x
 2017/05/05 10:45:26
```

2. Specify the numeric slot value to export the tape from the DD VTL.

   Run the following command: `vtl export cloud-vtl-pool slot 1 count 1`

3. Verify the tape is in the vault.

   Run the following command: `vtl tape show vault`

4. Select the tape for data movement.

   Run the following command: `vtl tape select-for-move barcode T00001L3 count 1 pool cloud-vtl-pool to-tier cloud`

   NOTE: If the data movement policy is age-threshold, data movement occurs automatically after 15-20 minutes.

5. View the list of tapes scheduled to move to cloud storage during the next data movement operation. The tapes selected for movement display an `(S)` in the location column.

   Run the following command: `vtl tape show vault`

```
Processing tapes.....
Barcode    Pool               Location    State      Size          Used (%)    Comp
Modification Time
--------   ----------------   ---------   ------   ------   ----------------   ----
-------------------
T00003L3   cloud-vtl-pool     vault (S)   RW        5 GiB   5.0 GiB (99.07%)    63x
2017/05/05 10:43:43
T00006L3   cloud-vtl-pool     ecs-unit1   n/a       5 GiB   5.0 GiB (99.07%)    62x
2017/05/05 10:45:49
--------   ----------------   ---------   ------   ------   ----------------   ----
-------------------
* RD : Replication Destination
(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next
data-movement run.
(R) Recall operation is in progress for the tape.
```

```
VTL Tape Summary
----------------
Total number of tapes:      4024
Total pools:                3
Total size of tapes:        40175 GiB
Total space used by tapes: 39.6 GiB
Average Compression:        9.7x
```

6. If the data movement policy is user-managed, initiate the data movement operation.

   Run the following command: `data-movement start`

7. Observe the status of the data movement operation.

   Run the following command: `data-movement watch`

8. Verify the tape volumes successfully move to cloud storage.

   Run the following command: `vtl tape show all cloud-unit ecs-unit1`

```
Processing tapes.....
Barcode  Pool           Location  State  Size         Used (%) Comp Modification Time
-------- -------------- --------- ----- ----- ---------------- ---- -------------------
T00001L3 cloud-vtl-pool ecs-unit1 n/a   5 GiB 5.0 GiB (99.07%)  89x 2017/05/05 10:41:41
T00006L3 cloud-vtl-pool ecs-unit1 n/a   5 GiB 5.0 GiB (99.07%)  62x 2017/05/05 10:45:49
-------- -------------- --------- ----- ----- ---------------- ---- -------------------
(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next
data-movement run.
(R) Recall operation is in progress for the tape.

VTL Tape Summary
----------------
Total number of tapes:      4
Total pools:                2
Total size of tapes:        16 GiB
Total space used by tapes: 14.9 GiB
Average Compression:        59.5x
```

Restore data held in the cloud

When a client requests data for restore from the backup application server, the backup application should generate an alert or message requesting the required volumes from the cloud unit.

The volume must be recalled from the cloud and checked into the DD VTL library before the backup application must be notified of the presence of the volumes.

NOTE: Verify the backup application is aware of status changes for volumes moved to cloud storage. Complete the necessary steps for the backup application to refresh its inventory to reflect the latest volume status.

Manually recall a tape volume from cloud storage

Recall a tape from Cloud Tier to the local VTL vault.

Steps

1. Select **Protocols** > **DD VTL**.
2. Expand the list of pools, and select the pool which is configured to migrate tapes to Cloud Tier.
3. In the pool pane, click the **Tape** tab.
4. Select one or more tapes that are located in a cloud unit.
5. Click **Recall Cloud Tapes** to recall tapes from Cloud Tier.

Results

After the next scheduled data migration, the tapes are recalled from the cloud unit to the vault. From the vault, the tapes can be returned to a library.

CLI equivalent

Steps

1. Identify the volume required to restore data.
2. Recall the tape volume from the vault.
   Run the following command: `vtl tape recall start barcode T00001L3 count 1 pool cloud-vtl-pool`
3. Verify the recall operation started.
   Run the following command: `data-movement status`
4. Verify the recall operation completed successfully.
   Run the following command: `vtl tape show all barcode T00001L3`

```
Processing tapes....
Barcode    Pool             Location           State    Size           Used (%)    Comp
Modification Time
--------   --------------   ----------------   -----    -----   ----------------   ----
-------------------
T00001L3   cloud-vtl-pool   cloud-vtl slot 1   RW       5 GiB   5.0 GiB (99.07%)   239x
2017/05/05 10:41:41
--------   --------------   ----------------   -----    -----   ----------------   ----
-------------------
(S) Tape selected for migration to cloud. Selected tapes will move to cloud on the next
data-movement run.
(R) Recall operation is in progress for the tape.

VTL Tape Summary
----------------
Total number of tapes:     1
Total pools:               1
Total size of tapes:       5 GiB
Total space used by tapes: 5.0 GiB
Average Compression:       239.1x
```

5. Validate the file location.
   Run the following command: `filesys report generate file-location path /data/col1/cloud-vtl-pool`

```
filesys report generate file-location path /data/col1/cloud-vtl-pool
--------------------------------      --------------------------
File Name                             Location(Unit Name)
--------------------------------      --------------------------
/data/col1/cloud-vtl-pool/.vtl_pool     Active
/data/col1/cloud-vtl-pool/.vtc/T00001L3 Active
-----------------------------------------------------------------
```

6. Import the recalled tape to the DD VTL.
   Run the following command: `vtl import cloud-vtl barcode T00001L3 count 1 pool cloud-vtl-pool element slot`

```
imported 1 tape(s)...sysadmin@ddbeta70# vtl tape show cloud-vtlProcessing tapes.....
```

7. Check the volume into the backup application inventory.
8. Restore data through the backup application.
9. When restore is completed check the tape volume out of the backup application inventory.
10. Export the tape volume from the DD VTL to the vault.
11. Move the tape back to the cloud unit.

Working with access groups

Access groups hold a collection of initiator WWPNs (worldwide port names) or aliases and the drives and changers they are allowed to access. A DD VTL default group named *TapeServer* lets you add devices that will support NDMP (Network Data Management Protocol)-based backup applications.

Access group configuration allows initiators (in general backup applications) to read and write data to devices in the same access group.

Access groups let clients access only selected LUNs (media changers or virtual tape drives) on a system. A client set up for an access group can access only devices in its access group.

Avoid making access group changes on a DD system during active backup or restore jobs. A change may cause an active job to fail. The impact of changes during active jobs depends on a combination of backup software and host configurations.

Selecting **Access Groups** > **Groups** displays the following information for all access groups.

Table 67. Access group information

Item	Description
Group Name	Name of group.
Initiators	Number of initiators in group.
Devices	Number of devices in group.

If you select **View All Access Groups**, you are taken to the Fibre Channel view.

From the More Tasks menu, you can create or delete a group.

Related concepts

Fibre Channel view
Managing a DD VTL

Related tasks

Verifying connectivity and creating access groups
Configuring the NDMP device TapeServer group

Creating an access group

Access groups manage access between devices and initiators. Do not use the default TapeServer access group unless you are using NDMP.

Steps

1. Select **Access Groups** > **Groups**.
2. Select **More Tasks** > **Group** > **Create**
3. In the Create Access Group dialog, enter a name, from 1 to 128 characters, and select **Next.**
4. Add devices, and select **Next.**
5. Review the summary, and select **Finish** or **Back**, as appropriate.

   CLI Equivalent

```
# vtl group create My_Group
```

Adding an access group device

Access group configuration allows initiators (in general backup applications) to read and write data to devices in the same access group.

Steps

1. Select **Access Groups** > **Groups**. You can also select a specific *group*.
2. Select **More Tasks** > **Group** > **Create** or **Group** > **Configure**.
3. In the Create or Modify Access Group dialog, enter or modify the **Group Name** if desired. (This field is required.)
4. To configure initiators to the access group, check the box next to the initiator. You can add initiators to the group later.
5. Select **Next**.
6. In the Devices display, select Add (**+**) to display the Add Devices dialog.

a. Verify that the correct library is selected in the Library Name drop-down list, or select another library.
b. In the Device area, select the checkboxes of the devices (changer and drives) to be included in the group.
c. Optionally, specify a starting LUN in the LUN Start Address text box.

   This is the LUN that the DD system returns to the initiator. Each device is uniquely identified by the library and the device name. (For example, it is possible to have drive 1 in Library 1 and drive 1 in Library 2). Therefore, a LUN is associated with a device, which is identified by its library and device name.

   When presenting LUNs via attached FC ports on FC HBA/SLIC, ports can be designated as primary, secondary, or none. A Primary port for a set of LUNs is the port that is currently advertizing those LUNs to a fabric. A secondary port is a port that will broadcast a set of LUNs in the event of primary path failure (this requires manual intervention). A setting of none is used in the case where you do not wish to advertize selected LUNs. The presentation of LUNs depends on the SAN topology in question.

   The initiators in the access group interact with the LUN devices that are added to the group.

   The maximum LUN accepted when creating an access group is 16383.

   A LUN can be used only once for an individual group. The same LUN can be used with multiple groups.

   Some initiators (clients) have specific rules for target LUN numbering; for example, requiring LUN 0 or requiring contiguous LUNs. If these rules are not followed, an initiator may not be able to access some or all of the LUNs assigned to a DD VTL target port.

   Check your initiator documentation for special rules, and if necessary, alter the device LUNs on the DD VTL target port to follow the rules. For example, if an initiator requires LUN 0 to be assigned on the DD VTL target port, check the LUNs for devices assigned to ports, and if there is no device assigned to LUN 0, change the LUN of a device so it is assigned to LUN 0.

d. In the Primary and Secondary Endpoints area, select an option to determine from which ports the selected device will be seen. The following conditions apply for designated ports:
   - all – The checked device is seen from all ports.
   - none – The checked device is not seen from any port.
   - select – The checked device is to be seen from selected ports. Select the checkboxes of the appropriate ports.

     If only primary ports are selected, the checked device is visible only from primary ports.

     If only secondary ports are selected, the checked device is visible only from secondary ports. Secondary ports can be used if the primary ports become unavailable.

   The switchover to a secondary port is not an automatic operation. You must manually switch the DD VTL device to the secondary ports if the primary ports become unavailable.

   The port list is a list of physical port numbers. A port number denotes the PCI slot and a letter denotes the port on a PCI card. Examples are 1a, 1b, or 2a, 2b.

   A drive appears with the same LUN on all the ports that you have configured.

e. Select **OK**.

   You are returned to the Devices dialog box where the new group is listed. To add more devices, repeat these five substeps.

7. Select **Next**.
8. Select **Close** when the `Completed` status message is displayed.

CLI Equivalent

```
# vtl group add VTL_Group vtl NewVTL changer lun 0 primary-port all secondary-port all
# vtl group add VTL_Group vtl NewVTL drive 1 lun 1 primary-port all secondary-port all
# vtl group add SetUp_Test vtl SetUp_Test drive 3 lun 3 primary-port endpoint-fc-0
secondary-port endpoint-fc-1

# vtl group show Setup_Test
Group: SetUp_Test

Initiators:
Initiator Alias   Initiator WWPN
---------------   ----------------------
tsm6_p23          21:00:00:24:ff:31:ce:f8
---------------   ----------------------
```

```
Devices:
Device Name          LUN   Primary Ports   Secondary Ports    In-use Ports
------------------   ---   -------------   ---------------   -------------
SetUp_Test changer     0             all               all             all
SetUp_Test drive 1     1             all               all             all
SetUp_Test drive 2     2              5a                5b              5a
SetUp_Test drive 3     3   endpoint-fc-0     endpoint-fc-1   endpoint-fc-0
------------------   ---   -------------   ---------------   -------------
```

Related concepts

Working with initiators
Working with resources

Related tasks

Verifying connectivity and creating access groups
Configuring the NDMP device TapeServer group

Modifying or deleting an access group device

You may need to modify or delete a device from an access group.

Steps

1. Select **Protocols** > **VTL** > **Access Groups** > **Groups** > *group*.
2. Select **More Tasks** > **Group** > **Configure**.
3. In the Modify Access Group dialog, enter or modify the **Group Name**. (This field is required.)
4. To configure initiators to the access group, check the box next to the initiator. You can add initiators to the group later.
5. Select **Next**.
6. Select a device, and select the edit (pencil) icon to display the Modify Devices dialog. Then, follow steps a-e. If you simply want to delete the device, select the delete (X) icon, and skip to step e.
   a. Verify that the correct library is selected in the Library drop-down list, or select another library.
   b. In the Devices to Modify area, select the checkboxes of the devices (Changer and drives) to be modified.
   c. Optionally, modify the starting LUN (logical unit number) in the LUN Start Address box.

      This is the LUN that the DD system returns to the initiator. Each device is uniquely identified by the library and the device name. (For example, it is possible to have drive 1 in Library 1 and drive 1 in Library 2). Therefore, a LUN is associated with a device, which is identified by its library and device name.

      The initiators in the access group interact with the LUN devices that are added to the group.

      The maximum LUN accepted when creating an access group is 16383.

      A LUN can be used only once for an individual group. The same LUN can be used with multiple groups.

      Some initiators (clients) have specific rules for target LUN numbering; for example, requiring LUN 0 or requiring contiguous LUNs. If these rules are not followed, an initiator may not be able to access some or all of the LUNs assigned to a DD VTL target port.

      Check your initiator documentation for special rules, and if necessary, alter the device LUNs on the DD VTL target port to follow the rules. For example, if an initiator requires LUN 0 to be assigned on the DD VTL target port, check the LUNs for devices assigned to ports, and if there is no device assigned to LUN 0, change the LUN of a device so it is assigned to LUN 0.

   d. In the Primary and Secondary Ports area, change the option that determines the ports from which the selected device is seen. The following conditions apply for designated ports:
      - all – The checked device is seen from all ports.
      - none – The checked device is not seen from any port.
      - select – The checked device is seen from selected ports. Select the checkboxes of the ports from which it will be seen.

        If only primary ports are selected, the checked device is visible only from primary ports.

If only secondary ports are selected, the checked device is visible only from secondary ports. Secondary ports can be used if primary ports become unavailable.

The switchover to a secondary port is not an automatic operation. You must manually switch the DD VTL device to the secondary ports if the primary ports become unavailable.

The port list is a list of physical port numbers. A port number denotes the PCI slot, and a letter denotes the port on a PCI card. Examples are 1a, 1b, or 2a, 2b.

A drive appears with the same LUN on all ports that you have configured.

e. Select **OK**.

Related concepts

Working with resources

Related tasks

Deleting an access group
Verifying connectivity and creating access groups

Deleting an access group

Before you can delete an access group, you must remove all of its initiators and LUNs.

Steps

1. Remove all of the initiators and LUNs from the group.
2. Select **Access Groups** > **Groups**.
3. Select **More Tasks** > **Group** > **Delete**.
4. In the Delete Group dialog, select the checkbox of the group to be removed, and select **Next**.
5. In the groups confirmation dialog, verify the deletion, and select **Submit**.
6. Select **Close** when the Delete Groups Status displays `Completed`.

CLI Equivalent

```
# scsitarget group destroy My_Group
```

Related tasks

Modifying or deleting an access group device

Working with a selected access group

Selecting **Access Groups** > **Groups** > *group* displays the following information for a selected access group.

Table 68. LUNs tab

Item	Description
LUN	Device address – maximum number is 16383. A LUN can be used only once within a group, but can be used again within another group. DD VTL devices added to a group must use contiguous LUNs.
Library	Name of library associated with LUN.
Device	Changers and drives.
In-Use Endpoints	Set of endpoints currently being used: primary or secondary.
Primary Endpoints	Initial (or default) endpoint used by backup application. In the event of a failure on this endpoint, the secondary endpoints may be used, if available.

Table 68. LUNs tab (continued)

Item	Description
Secondary Endpoints	Set of fail-over endpoints to use if primary endpoint fails.

Table 69. Initiators tab

Item	Description
Name	Name of initiator, which is either the WWPN or the alias assigned to the initiator.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel port.

From the More Tasks menu, with a group selected, you can configure that group, or set endpoints in use.

Related tasks

Verifying connectivity and creating access groups
Deleting access groups
Selecting endpoints for a device

Selecting endpoints for a device

Since endpoints connect a device to an initiator, use this process to set up the endpoints before you connect the device.

Steps

1. Select **Access Groups** > **Groups** > *group*.
2. Select **More Tasks** > **Endpoints** > **Set In-Use**.
3. In the Set in-Use Endpoints dialog, select only specific devices, or select **Devices** to select all devices in the list.
4. Indicate whether the endpoints are primary or secondary.
5. Select **OK**.

Related tasks

Verifying connectivity and creating access groups
Deleting access groups

Configuring the NDMP device TapeServer group

The DD VTL TapeServer group holds tape drives that interface with NDMP (Network Data Management Protocol)-based backup applications and that send control information and data streams over IP (Internet Protocol) instead of Fibre Channel (FC). A device used by the NDMP TapeServer must be in the DD VTL group TapeServer and is available *only* to the NDMP TapeServer.

Steps

1. Add tape drives to a new or existing library (in this example, named "dd9900-16").
2. Create slots and CAPs for the library.
3. Add the created devices in a library (in this example, "dd9900-16") to the TapeServer access group.
4. Enable the NDMP daemon by entering at the command line:

```
# ndmpd enable
Starting NDMP daemon, please wait.............
NDMP daemon is enabled.
```

5. Ensure that the NDMP daemon sees the devices in the TapeServer group:

```
# ndmpd show devicenames
NDMP Device          Virtual Name       Vendor   Product       Serial Number
-----------------    ----------------   ------   -----------   -------------
```

```
/dev/dd_ch_c0t0l0   dd9900-16 changer  STK      L180          6290820000
/dev/dd_st_c0t1l0   dd9900-16 drive 1  IBM      ULTRIUM-TD3   6290820001
/dev/dd_st_c0t2l0   dd9900-16 drive 2  IBM      ULTRIUM-TD3   6290820002
/dev/dd_st_c0t3l0   dd9900-16 drive 3  IBM      ULTRIUM-TD3   6290820003
/dev/dd_st_c0t4l0   dd9900-16 drive 4  IBM      ULTRIUM-TD3   6290820004
-----------------   ----------------   ------   -----------   -------------
```

6. Add an NDMP user (ndmp in this example) with the following command:

```
# ndmpd user add ndmp
Enter password:
Verify password:
```

7. Verify that user ndmp is added correctly:

```
# ndmpd user show
ndmp
```

8. Display the NDMP configuration:

```
# ndmpd option show all
Name              Value
--------------    --------
authentication    text
debug             disabled
port              10000
preferred-ip
--------------    --------
```

9. Change the default user password authentication to use MD5 encryption for enhanced security, and verify the change (notice the authentication value changed from text to md5):

```
# ndmpd option set authentication md5
# ndmpd option show all
Name              Value
--------------    --------
authentication    md5
debug             disabled
port              10000
preferred-ip
--------------    --------
```

Results

NDMP is now configured, and the TapeServer access group shows the device configuration. See the `ndmpd` chapter of the *DDOS Command Reference Guide* for the complete command set and options.

Related tasks

Creating drives
Adding slots
Adding CAPs
Adding an access group device
Verifying connectivity and creating access groups
Deleting access groups

Working with resources

Selecting **Resources** > **Resources** displays information about initiators and endpoints. An *initiator* is a backup client that connects to a system to read and write data using the Fibre Channel (FC) protocol. A specific initiator can support DD Boost over FC or DD VTL, but not both. An *endpoint* is the logical target on a DD system to which the initiator connects.

Table 70. Initiators tab

Item	Description
Name	Name of initiator, which is either the WWPN or the alias assigned to the initiator.

Table 70. Initiators tab (continued)

Item	Description
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
Online Endpoints	Group name where ports are seen by initiator. Displays `None` or `Offline` if the initiator is unavailable.

Table 71. Endpoints tab

Item	Description
Name	Specific name of endpoint.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
System Address	System address for the endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either `Yes` (enabled) or `No` (not enabled).
Status	DD VTL link status, which is either `Online` (capable of handling traffic) or `Offline`.

Configure Resources

Selecting **Configure Resources** takes you to the Fibre Channel area, where you can configure endpoints and initiators.

Related concepts

Fibre Channel view
Managing a DD VTL

Working with initiators

Selecting **Resources** > **Resources** > **Initiators** displays information about initiators. An *initiator* is a client system FC HBA (fibre channel host bus adapter) WWPN (worldwide port name) with which the DD system interfaces. An *initiator name* is an alias for the client's WWPN, for ease of use.

While a client is mapped as an initiator – but before an access group has been added – the client cannot access any data on a DD system.

After adding an access group for the initiator or client, the client can access only the devices in that access group. A client can have access groups for multiple devices.

An access group may contain multiple initiators, but an initiator can exist in only one access group.

NOTE: A maximum of 1024 initiators can be configured for a DD system.

Table 72. Initiator information

Item	Description
Name	Name of initiator.
Group	Group associated with initiator.
Online Endpoints	Endpoints seen by initiator. Displays `none` or `offline` if initiator is unavailable.

Table 72. Initiator information (continued)

Item	Description
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
Vendor Name	Name of vendor for initiator.

Selecting **Configure Initiators** takes you to the Fibre Channel area, where you can configure endpoints and initiators.

CLI Equivalent

```
# scsitarget initiator show detailed
Initiator:                NewAliasName
System Address:           21:00:f4:e9:d4:59:d6:30
Group:                    DDBOOST_DFC_Group
Service:                  DD-Boost FC
Address Method:           auto
Transport:                FibreChannel
FC WWPN:                  21:00:f4:e9:d4:59:d6:30
FC WWNN:                  20:00:f4:e9:d4:59:d6:30
FC Symbolic Port Name:    QLE2692 FW:v8.08.204 DVR:v10.00.00.12.07.7-k

Initiator      Status    Endpoint
------------   ------    -------------
NewAliasName   Online    endpoint-fc-0
------------   ------    -------------
```

Related concepts

Fibre Channel view

Working with endpoints

Selecting **Resources** > **Resources** > **Endpoints** provides information about endpoint hardware and connectivity.

Table 73. Hardware Tab

Item	Description
System Address	System address of endpoint.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
Enabled	HBA (host bus adapter) port operational state, which is either `Yes` (enabled) or `No` (not enabled).
NPIV	NPIV status of this endpoint: eithe Enabled or Disabled.
Link Status	Link status of this endpoint: either Online or Offline.
Operation Status	Operation status of this endpoint: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this endpoint.

Table 74. Endpoints Tab

Item	Description
Name	Specific name of endpoint.

Table 74. Endpoints Tab (continued)

Item	Description
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel (FC) port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
System Address	System address of endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either `Yes` (enabled) or `No` (not enabled).
Link Status	Link status of this endpoint: either Online or Offline.

Configure Endpoints

Selecting **Configure Endpoints** takes you to the Fibre Channel area, where you can change any of the above information for the endpoint.

CLI Equivalent

```
# scsitarget endpoint show list
Endpoint       System Address  Transport     Enabled  Status
--------       --------------  ---------     -------  ------
endpoint-fc-0  5a              FibreChannel  Yes      Online
endpoint-fc-1  5b              FibreChannel  Yes      Online
```

Working with a selected endpoint

Selecting **Resources** > **Resources** > **Endpoints** > *endpoint* provides information about the endpoint's hardware, connectivity, and statistics.

Table 75. Hardware tab

Item	Description
System Address	System address of endpoint.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.
Enabled	HBA (host bus adapter) port operational state, which is either `Yes` (enabled) or `No` (not enabled).
NPIV	NPIV status of this endpoint: eithe Enabled or Disabled.
LInk Status	Link status of this endpoint: either Online or Offline.
Operation Status	Operation status of this endpoint: either Normal or Marginal.
# of Endpoints	Number of endpoints associated with this endpoint.

Table 76. Summary tab

Item	Description
Name	Specific name of endpoint.
WWPN	Unique worldwide port name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the Fibre Channel port.
WWNN	Unique worldwide node name, which is a 64-bit identifier (a 60-bit value preceded by a 4-bit *Network Address Authority* identifier), of the FC node.

Table 76. Summary tab (continued)

Item	Description
System Address	System address of endpoint.
Enabled	HBA (host bus adapter) port operational state, which is either `Yes` (enabled) or `No` (not enabled).
LInk Status	Link status of this endpoint: either Online or Offline.

Table 77. Statistics tab

Item	Description
Endpoint	Specific name of endpoint.
Library	Name of library containing the endpoint.
Device	Number of device.
Ops/s	Operations per second.
Read KiB/s	Speed of reads in KiB per second.
Write KiB/s	Speed of writes in KiB per second.

Table 78. Detailed Statistics tab

Item	Description
Endpoint	Specific name of endpoint.
# of Control Commands	Number of control commands.
# of Read Commands	Number of read commands.
# of Write Commands	Number of write commands.
In (MiB)	Number of MiB written (the binary equivalent of MB).
Out (MiB)	Number of MiB read.
# of Error Protocol	Number of error protocols.
# of Link Fail	Number of link failures.
# of Invalid Crc	Number of invalid CRCs (cyclic redundancy checks).
# of Invalid TxWord	Number of invalid tx (transmission) words.
# of Lip	Number of LIPs (loop initialization primitives).
# of Loss Signal	Number of signals or connections that have been lost.
# of Loss Sync	Number of signals or connections that have lost synchronization.

Working with pools

Selecting **Pools** > **Pools** displays detailed information for the Default pool and any other existing pools. A *pool* is a collection of tapes that maps to a directory on the file system. Pools are used to replicate tapes to a destination. You can convert directory-based pools to MTree-based pools to take advantage of the greater functionality of MTrees.

Note the following about pools:

- Pools can be of two types: MTree (recommended), or Directory, which is backward-compatible.
- A pool can be replicated no matter where individual tapes are located. Tapes can be in the vault or in a library (slot, cap, or drive).
- You can copy and move tapes from one pool to another.
- Pools are not accessible by backup software.
- No DD VTL configuration or license is needed on a replication destination when replicating pools.
- You must create tapes with unique barcodes. Duplicate barcodes may cause unpredictable behavior in backup applications and can be confusing to users.

- Two tapes in two different pools on a DD system may have the same name, and in this case, neither tape can be moved to the other tape's pool. Likewise, a pool sent to a replication destination must have a name that is unique on the destination.

Table 79. Pools tab

Item	Description
Name	The name of the pool.
Type	Whether it is a Directory or MTree pool.
Status	The status of the pool.
Tape Count	The number of tapes in the pool.
Size	The total configured data capacity of tapes in the pool, in GiB (Gibibytes base-2 equivalent of GB, Gigabytes).
Physical Used	The amount of space used on virtual tapes in the pool.
Compression	The average amount of compression achieved for data on tapes in the pool.
Cloud Unit	The name of the cloud unit where the DD VTL pool migrates data.
Cloud Data Movement Policy	The data movement policy that governs migration of DD VTL data to Cloud Tier storage.

Table 80. Replication tab

Item	Description
Name	The name of the pool.
Configured	Whether replication is configured for the pool: yes or no.
Remote Source	Contains an entry only if the pool is replicated from another DD system.
Remote Destination	Contains an entry only if the pool replicates to another DD system.

From the More Tasks menu, you can create and delete pools, as well as search for tapes.

Related concepts

Managing a DD VTL

Related tasks

Deleting pools
Searching for tapes

Creating pools

You can create backward-compatible pools, if necessary for your setup.

Steps

1. Select **Pools** > **Pools**.
2. Select **More Tasks** > **Pool** > **Create**.
3. In the Create Pool dialog, enter a Pool Name, noting that a pool name:
   - cannot be "all," "vault," or "summary."
   - cannot have a space or period at its beginning or end.
   - is case-sensitive.
4. If you want to create a directory pool (which is backward compatible with the previous version of DD System Manager), select the option "Create a directory backwards compatibility mode pool. " However, be aware that the advantages of using an MTree pool include the ability to:
   - make individual snapshots and schedule snapshots.
   - apply retention locks.

- set an individual retention policy.
- get compression information.
- get data migration policies to the Retention Tier.
- establish a storage space usage policy (quota support) by setting hard limits and soft limits.

5. Select **OK** to display the Create Pool Status dialog.
6. When the Create Pool Status dialog shows `Completed`, select **Close**. The pool is added to the Pools subtree, and you can now add virtual tapes to it.

CLI Equivalent

```
# vtl pool add VTL_Pool
A VTL pool named VTL_Pool is added.
```

Deleting pools

Before a pool can be deleted, you must have deleted any tapes contained within it. If replication is configured for the pool, the replication pair must also be deleted. Deleting a pool corresponds to renaming the MTree and then deleting it, which occurs at the next cleaning process.

Steps

1. Select **Pools** > **Pools** > ***pool***.
2. Select **More Tasks** > **Pool** > **Delete**.
3. In the Delete Pools dialog, select the checkbox of items to delete:
   - The name of each pool, or
   - **Pool Names**, to delete all pools.
4. Select **Submit** in the confirmation dialogs.
5. When the Delete Pool Status dialog shows `Completed`, select **Close**.

   The pool will have been removed from the Pools subtree.

Working with a selected pool

Both **Virtual Tape Libraries** > **VTL Service** > **Vault** > *pool* and **Pools** > **Pools** > *pool* display detailed information for a selected pool. Notice that pool "Default" always exists.

Pool tab

Table 81. Summary

Item	Description
Convert to MTree Pool	Select this button to convert a Directory pool to an MTree pool.
Type	Whether it is a Directory or MTree pool.
Tape Count	The number of tapes in the pool.
Capacity	The total configured data capacity of tapes in the pool, in GiB (Gibibytes, base-2 equivalent of GB, Gigabytes).
Logical Used	The amount of space used on virtual tapes in the pool.
Compression	The average amount of compression achieved for data on tapes in the pool.

Table 82. Pool Tab: Cloud Data Movement - Protection Distribution

Item	Description
Pool type (%)	VTL Pool and Cloud (if applicable), with the current percentage of data in parentheses.
Name	Name of the local VTL pool, or cloud provider.
Logical Used	The amount of space used on virtual tapes in the pool.
Tape Count	The number of tapes in the pool.

Table 83. Pool Tab: Cloud Data Movement - Cloud Data Movement Policy

Item	Description
Policy	Age of tapes in days, or manual selection.
Older Than	Age threshold for an age-based data movement policy.
Cloud Unit	Destination cloud unit.

Tape tab

Table 84. Tape controls

Item	Description
Create	Create a new tape.
Delete	Delete the selected tapes.
Copy	Make a copy of a tape.
Move between Pool	Move the selected tapes to a different pool.
Select for Cloud Move[a]	Schedule the selected tapes for migration to Cloud Tier.
Unselect from Cloud Move[a]	Remove the selected tapes from the schedule for migration to Cloud Tier.
Recall Cloud Tapes	Recall the selected tapes from Cloud Tier.
Move to Cloud Now	Migrate the selected tapes to Cloud Tier without waiting for the next scheduled migration.

a. This option is only available if the data movement policy is configured for manual selection.

Table 85. Tape information

Item	Description
Barcode	Tape barcode.
Size	Maximum size of the tape.
Physical Used	Physical storage capacity used by the tape.
Compression	Compression ratio on the tape.
Location	Location of the tape.
Modification Time	Last time the tape was modified.
Recall Time	Last time the tape was recalled.

Replication tab

Table 86. Replication

Item	Description
Name	The name of the pool.
Configured	Whether replication is configured for this pool: yes or no.
Remote Source	Contains an entry only if the pool is replicated from another DD system.
Remote Destination	Contains an entry only if the pool replicates to another DD system.

You can also select the **Replication Detail** button, at the top right, to go directly to the Replication information panel for the selected pool.

From either the Virtual Tape Libraries or Pools area, from the More Tasks menu, you can create, delete, move, copy, or search for a tape in the pool.

From the Pools area, from the More Tasks menu, you can rename or delete a pool.

Related tasks

Creating tapes
Deleting tapes
Searching for tapes
Creating an MTree or pool replication pair
Viewing estimated completion time for backup jobs

Converting a directory pool to an MTree pool

MTree pools have many advantages over directory pools. See the *Creating pools* section for more information.

Steps

1. Make sure the following prerequisites have been met:
   - The source and destination pools must have been synchronized, so that the number of tapes, and the data on each side, remains intact.
   - The directory pool must not be a replication source or destination.
   - The file system must not be full.
   - The file system must not have reached the maximum number of MTrees allowed.
   - There must not already be an MTree with the same name.
2. With the directory pool you wish to convert highlighted, choose **Convert to MTree Pool**.
3. Select **OK** in the Convert to MTree Pool dialog.
4. Be aware that conversion affects replication in the following ways:
   - DD VTL is temporarily disabled on the replicated systems during conversion.
   - The destination data is copied to a new pool on the destination system to preserve the data until the new replication is initialized and synced. Afterward, you may safely delete this temporarily copied pool, which is named **CONVERTED-***pool*, where *pool* is the name of the pool that was upgraded (or the first 18 characters for long pool names).
   - DD Retention Lock cannot be enabled on systems involved in MTree pool conversion.

Moving tapes between pools

If they reside in the vault, tapes can be moved between pools to accommodate replication activities. For example, pools are needed if all tapes were created in the Default pool, but you later need independent groups for replicating groups of tapes. You can create named pools and re-organize the groups of tapes into new pools.

Steps

1. With a pool highlighted, select **More Tasks** > **Tapes** > **Move**.

Note that when started from a pool, the Tapes Panel allows tapes to be moved only between pools.

2. In the Move Tapes dialog, enter information to search for the tapes to move, and select **Search**:

Table 87. Move Tapes dialog

Field	User input
Location	Location cannot be changed.
Pool	Select the name of the pool where the tapes reside. If no pools have been created, use the Default pool.
Barcode	Specify a unique barcode. or leave the default (*) to import a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be returned to you. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

3. From the search results list, select the tapes to move.
4. From the Select Destination: Location list, select the location of the pool to which tapes are to be moved. This option is available only when started from the (named) Pool view.
5. Select **Next**.
6. From the Move Tapes view, verify the summary information and tape list, and select **Submit**.
7. Select **Close** in the status window.

Copying tapes between pools

Tapes can be copied between pools, or from the vault to a pool, to accommodate replication activities. This option is available only when started from the (named) Pool view.

Steps

1. With a pool highlighted, select **More Tasks** > **Tapes** > **Copy**.
2. In the Copy Tapes Between Pools dialog. select the checkboxes of tapes to copy, or enter information to search for the tapes to copy, and select **Search**:

Table 88. Copy Tapes Between Pools dialog

Field	User input
Location	Select either a library or the **Vault** for locating the tape. While tapes always show up in a pool (under the Pools menu), they are technically in either a library or the vault, but not both, and they are never in two libraries at the same time. Use the import/export options to move tapes between the vault and a library.
Pool	To copy tapes between pools, select the name of the pool where the tapes currently reside. If no pools have been created, use the **Default** pool.
Barcode	Specify a unique barcode. or leave the default (*) to import a group of tapes. Barcode allows the wildcards ? and *, where ? matches any single character and * matches 0 or more characters.
Count	Enter the maximum number of tapes you want to be imported. If you leave this blank, the barcode default (*) is used.
Tapes Per Page	Select the maximum number of tapes to display per page. Possible values are 15, 30, and 45.
Items Selected	Shows the number of tapes selected across multiple pages – updated automatically for each tape selection.

3. From the search results list, select the tapes to copy.
4. From the Select Destination: Pool list, select the pool where tapes are to be copied. If a tape with a matching barcode already resides in the destination pool, an error is displayed, and the copy aborts.
5. Select **Next**.

6. From the Copy Tapes Between Pools dialog, verify the summary information and the tape list, and select **Submit**.
7. Select **Close** on the Copy Tapes Between Pools Status window.

Renaming pools

A pool can be renamed only if none of its tapes is in a library.

Steps

1. Select **Pools** > **Pools** > ***pool***.
2. Select **More Tasks** > **Pool** > **Rename**.
3. In the Rename Pool dialog, enter the new Pool Name, with the caveat that this name:
   - cannot be "all," "vault," or "summary."
   - cannot have a space or period at its beginning or end.
   - is case-sensitive.
4. Select **OK** to display the Rename Pool status dialog.
5. After the Rename Pool status dialog shows `Completed`, select **OK**.

   The pool will have been renamed in the Pools subtree in both the Pools and the Virtual Tape Libraries areas.

23

DD Replicator

This chapter includes:

Topics:

DD Replicator overview

DD Replicator provides automated, policy-based, network-efficient, and encrypted replication for DR (disaster recovery) and multi-site backup and archive consolidation. DD Replicator asynchronously replicates only compressed, deduplicated data over a WAN (wide area network).

DD Replicator performs two levels of deduplication to significantly reduce bandwidth requirements: *local* and *cross-site* deduplication. Local deduplication determines the unique segments to be replicated over a WAN. Cross-site deduplication further reduces bandwidth requirements when multiple sites are replicating to the same destination system. With cross-site deduplication, any redundant segment previously transferred by any other site, or as a result of a local backup or archive, will not be replicated again. This improves network efficiency across all sites and reduces daily network bandwidth requirements up to 99%, making network-based replication fast, reliable, and cost-effective.

In order to meet a broad set of DR requirements, DD Replicator provides flexible replication topologies, such as full system mirroring, bi-directional, many-to-one, one-to-many, and cascaded. In addition, you can choose to replicate either all or a subset of the data on your DD system. For the highest level of security, DD Replicator can encrypt data being replicated between DD systems using the standard SSL (Secure Socket Layer) protocol.

DD Replicator scales performance and supported fan-in ratios to support large enterprise environments.

Before getting started with DD Replicator, note the following general requirements:

- DD Replicator is a licensed product. Contact your Dell sales representative to purchase licenses.
- You can usually replicate between machines that are within five releases of each other, for example, from 6.0 to 7.2. However, there may be exceptions to this, so review the tables in the *Replication version compatibility* section, or check with your Dell representative.
- If you are unable to manage and monitor DD Replicator from the current version of the DD System Manager, use the `replication` commands described in the *DDOS Command Reference Guide*.

Related concepts

Replication version compatibility

Prerequisites for replication configuration

Before configuring a replication, review the following prerequisites to minimize initial data transfer time, prevent overwriting of data, etc.

- **Contexts** – Determine the maximum number of contexts for your DD systems by reviewing the replication streams numbers in Best practices for data streams sent to DD systems .
- **Compatibility** – If you are using DD systems running different versions of DDOS, review the next section on Replication Version Compatibility.
- **Initial Replication** – If the source holds a lot of data, the initial replication operation can take many hours. Consider putting both DD systems in the same location with a high-speed, low-latency link. After the first replication, you can move the systems to their intended locations because only new data will be sent.
- **Bandwidth Delay Settings** – Both the source and destination must have the same bandwidth delay settings. These tuning controls benefit replication performance over higher latency links by controlling the TCP (transmission control protocol) buffer size. The source system can then send enough data to the destination while waiting for an acknowledgment.
- **Adequate Storage** – At a minimum, the destination must have the *same amount of space* as the source.
- **Security** – DDOS requires that port 3009 be open in order to configure secure replication over an Ethernet connection.

Replication version compatibility

To use DD systems running different versions of DDOS for a source or destination, the following tables provide compatibility information for single-node, DD Retention Lock, MTree, collection, delta (low bandwidth optimization), and cascaded replication.

In general:

- For DD Boost or DD Boost OST, see the *DD Boost for Partner Integration Administration Guide* or the *DD Boost for OpenStorage Administration Guide* for supported configurations.
- The recovery procedure is valid for all supported replication configurations.
- File migration is supported whenever collection replication is supported.
- For MTree replication or managed file replication, if a DDOS 7.2 or later source is configured to replicate to a target running DDOS 6.2 with gz or gzfast compression, the target system must be upgraded to DDOS 6.2.0.35 or higher.
- For collection replication, the source and destination systems must be on the same release. There is no compatibility between release families for collection replication.
- For cascaded configurations, the maximum number of hops is two, that is, three DD systems.
- One-to-many, many-to-one, and cascaded replication are supported according to the version compatibility listed in the table below.

NOTE: Directory replication is supported in DDOS 7.7.0.0 and lower, and not supported in DDOS versions 7.7.1.0 and later.

Replication compatibility (as either a source or destination system) with other versions of DDOS depends on the type of replication:

- MTree replication and MFR: Compatible with DDOS 6.2 and higher.
- Collection replication: Compatible only with this DDOS release family.
- Delta (low bandwidth optimization) replication: Compatible with the two previous DDOS release families and the next two DDOS release families.

  NOTE: For example, a system running DDOS 7.8.X can be the source or destination in a delta replication context with systems running DDOS 7.6.X, DDOS 7.7.X, DDOS 7.9.X, or DDOS 7.10.X.

TLS version support

By default, the system supports TLS versions 1.0, 1.1, and 1.2. However, it is possible to configure the system to support TLS version 1.2 only by changing the system parameter `REPL_SSL_DISABLE_TLSV1_0`. Changing the system parameter requires SE access to the system. Contact Dell Support if this change is required.

Replication types

Replication typically consists of a *source* DD system (which receives data from a backup system) and one or more *destination* DD systems. Each DD system can be the source and/or the destination for replication contexts. During replication, each DD system can perform normal backup and restore operations.

Each replication type establishes a *context* associated with an existing MTree on the source. The replicated context is created on the destination when a context is established. The context establishes a replication pair, which is always active, and any data landing in the source will be copied to the destination at the earliest opportunity. Paths configured in replication contexts are absolute references and do not change based on the system in which they are configured.

A protection system can be set up for collection or MTree replication.

- *Collection replication* duplicates the entire data store on the source and transfers that to the destination, and the replicated volume is read-only.
- *MTree replication* replicates entire MTrees (that is, a virtual file structure that enables advanced management). Media pools can also be replicated, and by default, an MTree is created for replication.

For any replication type, note the following requirements:

- A destination system must have available storage capacity that is at least the size of the expected maximum size of the source directory. Be sure that the destination system has enough network bandwidth and disk space to handle all traffic from replication sources.
- The file system must be enabled or, based on the replication type, will be enabled as part of the replication initialization.
- The source must exist.
- The destination must not exist.
- The destination will be created when a context is built and initialized.
- After replication is initialized, ownership and permissions of the destination are always identical to those of the source.
- In the replication command options, a specific replication pair is always identified by the destination.
- Both systems must have an active, visible route through the IP network so that each system can resolve its partner's host name.

The choice of replication type depends on your specific needs. The next sections provide descriptions and features of these two types, plus a brief introduction to Managed File Replication, which is used by DD Boost.

Managed file replication

Managed file replication, which is used by DD Boost, is a type of replication that is managed and controlled by backup software.

With managed file replication, backup images are directly transferred from one DD system to another, one at a time, at the request of the backup software.

The backup software keeps track of all copies, allowing easy monitoring of replication status and recovery from multiple copies.

Managed file replication offers flexible replication topologies including full system mirroring, bi-directional, many-to-one, one-to-many, and cascaded, enabling efficient cross-site deduplication.

Here are some additional points to consider about managed file replication:

- Replication contexts do not need to be configured.
- Lifecycle polices control replication of information with no intervention from the user.
- DD Boost will build and tear down contexts as needed on the fly.
- MFR automatically uses the most secure level of replication encryption configured on the replication pair.

For more information, see the `ddboost file-replication` commands in the *DDOS Command Reference Guide*.

MTree replication

MTree replication is used to replicate MTrees between DD systems. Periodic snapshots are created on the source, and the differences between them are transferred to the destination by leveraging the system's cross-site deduplication mechanism. This ensures that the data on the destination is always a point-in-time copy of the source, with file consistency. This also reduces replication of churn in the data, leading to more efficient utilization of the WAN.

The use of snapshots with MTree replication enables some intermediate changes to the source to be skipped. Skipping these changes further reduces the amount of data that is sent over the network, and therefore reduces replication lag.

With MTree replication, a DD system can be simultaneously the source of some replication contexts and the destination of other contexts. And that DD system can also receive data from backup and archive applications while it is replicating data.

MTree replication has the same flexible network deployment topologies and cross-site deduplication effects as managed file replication (the type used by DD Boost).

Here are some additional points to consider when using MTree replication:

- When replication is initialized, a destination read-only MTree is created automatically.
- Data can be logically separated into multiple MTrees to promote greater replication performance.
- Snapshots must be created on source contexts.
- Snapshots cannot be created on a replication destination.
- Snapshots are replicated with a fixed retention of one year; however, the retention is adjustable on the destination and must be adjusted there.
- Snapshots are not automatically deleted after breaking a replication context, and must be expired when they are no longer required to prevent the system from filling up. The following KB articles, available from the Online Support website, provide more information:
  - *Data Domain - Checking for Snapshots that are No Longer Needed* .
  - *Data Domain - Identifying Why a DDR is Filling Up* .
  - *Data Domain - Mtree_replication_resync_Snapshot_retention* .
- Replication contexts must be configured on both the source and the destination.
- Replicating DD VTL tape cartridges (or pools) simply means replicating MTrees or directories that contain DD VTL tape cartridges. Media pools are replicated by MTree replication, as a default. You cannot use the pool:// syntax to create replication contexts using the command line. When specifying pool-based replication in DD System Manager, MTree replication will be created.
- Replicating directories under an MTree is not permitted.
- A destination DD system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source MTree.
- After replication is initialized, ownership and permissions of the destination MTree are always identical to those of the source MTree. If the context is configured, the destination MTree is kept in a read-only state and can receive data only from the source MTree.
- At any time, due to differences in global compression, the source and destination MTree can differ in size.
- DD Retention Lock Compliance is supported with MTree replication, by default. If DD Retention Lock is licensed on a source, the destination must also have a DD Retention Lock license, or replication will fail. (To avoid this situation, you must disable DD Retention Lock.) If DD Retention Lock is enabled on a replication context, a replicated destination context will always contain data that is retention locked.
- DD Boost users should have the same user ID (UID) and primary group ID (GID) on both the source and destination systems.

MTree replication details

MTree replication involves the following steps:

1. A snapshot is created on the source replication context.
2. This snapshot is compared to the last previous snapshot.
3. Any differences between the two snapshots are sent to the destination replication context.
4. On the destination, the MTree is updated but no files are exposed to the user until all changes are received by the destination system.

These steps are repeated any time a snapshot is created on the source MTree. The following situations trigger the creation of a snapshot on the source system:

- System–generated periodic snapshot—When the replication lag is more than 15 minutes and there is no snapshot being currently replicated.
- User–created snapshot—At a time specified by the user, such as after the completion of a backup job.

For examples showing the interaction of different types of snapshots, see the KB article *How MTree Replication Works*, available from the Online Support website.

After the snapshot is replicated, the connection to the destination is closed. A new connection between the source and destination is established when the next snapshot is replicated.

Related concepts

MTrees overview

Collection replication

Collection replication performs whole-system mirroring in a one-to-one topology, continuously transferring changes in the underlying collection, including all of the logical directories and files of the DD file system.

Collection replication does not have the flexibility of the other types, but it can provide higher throughput and support more objects with less overhead, which may work better for high-scale enterprise cases.

Collection replication replicates the entire `/data/col1` area from a source DD system to a destination DD system.

NOTE: Collection replication is not supported for cloud-tier enabled systems.

Here are some additional points to consider when using collection replication:

- No granular replication control is possible. All data is copied from the source to the destination producing a read-only copy.
- Collection replication requires that the storage capacity of the destination system be equal to, or greater than, the capacity of the source system. If the destination capacity is less than the source capacity, the available capacity on the source is reduced to the capacity of the destination.
- The DD system to be used as the collection replication destination must be empty before configuring replication. After replication is configured, this system is dedicated to receive data from the source system.
- With collection replication, all user accounts and passwords are replicated from the source to the destination. However, as of DDOS 5.5.1.0, other elements of configuration and user settings of the DD system are not replicated to the destination; you must explicitly reconfigure them after recovery.
- Collection replication is supported with DD Secure Multitenancy (SMT). Core SMT information, contained in the registry namespace, including the tenant and tenant-unit definitions with matching UUIDs is automatically transferred during replication operation. However, the following SMT information is not automatically included for replication, and must be configured manually on the destination system:
  - Alert notification lists for each tenant-unit
  - All users assigned to the DD Boost protocol for use by SMT tenants, if DD Boost is configured on the system
  - The default-tenant-unit associated with each DD Boost user, if any, if DD Boost is configured on the system

  Using collection replication for disaster recovery with SMT describes how to manually configure these items on the replication destination.
- DD Retention Lock Compliance supports collection replication.
- Collection replication is not supported in cloud tier-enabled systems.
- With collection replication, data in a replication context on the source system that has not been replicated cannot be processed for file system cleaning. If file system cleaning cannot complete because the source and destination systems are out of sync, the system reports the cleaning operation status as `partial`, and only limited system statistics are available for the cleaning operation. If collection replication is disabled, the amount of data that cannot be processed for file system cleaning increases because the replication source and destination systems remain out of sync. The KB article *Data Domain: An overview of Data Domain File System (DDFS) clean/garbage collection (GC) phases*, available from the Online Support website provides additional information.
- To enhance throughput in a high bandwidth environment, run the **`replication modify <destination> crepl-gc-gw-optim`** command to disable collection replication bandwidth optimization.

Automatic Multi-Streaming (AMS)

Automatic Multi-Streaming (AMS) improves MTree replication performance. It uses multiple streams to replicate a single large file (32 GB or larger) to improve network bandwidth utilization during replication. By increasing the replication speed for individual files, AMS also improves the pipeline efficiency of the replication queue, and provides improved replication throughput and reduced replication lag.

When the workload presents multiple optimization choices, AMS automatically selects the best option for the workload. For example, if the workload is a large file with fastcopy attributes, the replication operation uses fastcopy optimization to avoid the overhead of scanning the file to identify unique segments between the replication pair. If the workload uses synthetics, replication uses synthetic replication on top of AMS to leverage local operations on the destination system for each replication stream to generate the file.

AMS is always enabled, and cannot be disabled.

Using DD Encryption with DD Replicator

DD Replicator can be used with the optional *DD Encryption* feature, enabling encrypted data to be replicated using collection or MTree replication

Replication contexts are always authenticated with a *shared secret*. That shared secret is used to establish a session key using a Diffie-Hellman key exchange protocol, and that session key is used to encrypt and decrypt the protection system encryption key when appropriate.

Each replication type works uniquely with encryption and offers the same level of security.

- *Collection replication* requires the source and destination to have the same encryption configuration, because the destination data is expected to be an exact replica of the source data. In particular, the encryption feature must be turned on or off at both the source and destination, and if the feature is turned on, the encryption algorithm and the system passphrases must also match. The parameters are checked during the replication association phase.

  During collection replication, the source transmits the data in encrypted form, and also transmits the encryption keys to the destination. The data can be recovered at the destination because the destination has the same passphrase and the same system encryption key.

  NOTE: Collection replication is not supported for cloud-tier enabled systems.
- *MTree replication* does not require encryption configuration to be the same at both the source and destination. Instead, the source and destination securely exchange the destination's encryption key during the replication association phase, and the data is re-encrypted at the source using the destination's encryption key before transmission to the destination.

  If the destination has a different encryption configuration, the data transmitted is prepared appropriately. For example, if the feature is turned off at the destination, the source decrypts the data, and it is sent to the destination un-encrypted.
- *Managed file replication (MFR)* automatically chooses the more secure encryption setting between the source and destination systems for each MFR job (per context). For example:
  - If the source system uses encryption with anonymous authentication and the destination system does not use encryption, MFR will use encryption with anonymous authentication.
  - If the source system uses encryption with one-way authentication and the destination system uses encryption with anonymous authentication, MFR will use encryption with one-way authentication.

  NOTE: If one of the source or destination system is running DDOS 7.10 or later and the other is running an older version of DDOS that does not support this automatic encryption selection for MFR, the encryption setting must be configured manually to make the source and destination systems match.
- In a *cascaded replication* topology, a replica is chained among three systems. The last system in the chain can be configured as a collection or MTree. If the last system is a collection replication destination, it uses the same encryption keys and encrypted data as its source. If the last system is an MTree replication destination, it uses its own key, and the data is encrypted at its source. The encryption key for the destination at each link is used for encryption. Encryption for systems in the chain works as in a replication pair.

Related concepts

DD Encryption overview

Replication topologies

DD Replicator supports five replication topologies (one-to-one, one-to-one bidirectional, one-to-many, many-to-one, and cascaded). The tables in this section show (1) how these topologies work with two types of replication (MTree and collection) and (2) how mixed topologies are supported with cascaded replication.

In general:

- Single node (SN) systems support all replication topologies.
- Single node-to-single node (SN -> SN) can be used for all replication types.
- Collection replication cannot be configured from either an SN system to a DD high availability-enabled system, nor from a DD high availability-enabled system to an SN system.
- For MTtree replication, DD high availability systems are treated like SN systems.
- Collection replication cannot be configured on Cloud Tier-enabled systems.

In this table:

- SN = single node DD system without Cloud Tier
- SN + CT = single node DD system with Cloud Tier

Table 89. Topology Support by Replication Type and DD System Type

Topologies	MTree Replication	Collection Replication
one-to-one	SN -> {SN | SN + CT}	SN -> SN
one-to-one bidirectional	SN -> {SN | SN + CT}	Not supported
one-to-many	SN -> {SN | SN + CT}	Not supported
many-to-one	SN -> {SN | SN + CT}	Not supported
cascaded	SN -> {SN | SN + CT} -> {SN | SN + CT}	Not supported

Cascaded replication supports mixed topologies where the second leg in a cascaded connection is different from the first type in a connection (for example, A -> B is MTree replication, and B -> C is collection replication).

Table 90. Mixed Topologies Supported with Cascaded Replication

Mixed Topologies	
SN – Dir Repl -> SN + CT – MTree Repl -> SN + CT – MTree Repl	SN – Dir Repl -> SN + CT – Col Repl -> SN + CT – Col Repl
SN – MTree Repl -> SN – Col Repl -> SN – Col Repl	SN – MTree Repl -> SN + CT – Col Repl -> SN + CT – Col Repl

One-to-one replication

The simplest type of replication is from a DD source system to a DD destination system, otherwise known as a *one-to-one* replication pair. This replication topology can be configured with MTree or collection replication types.

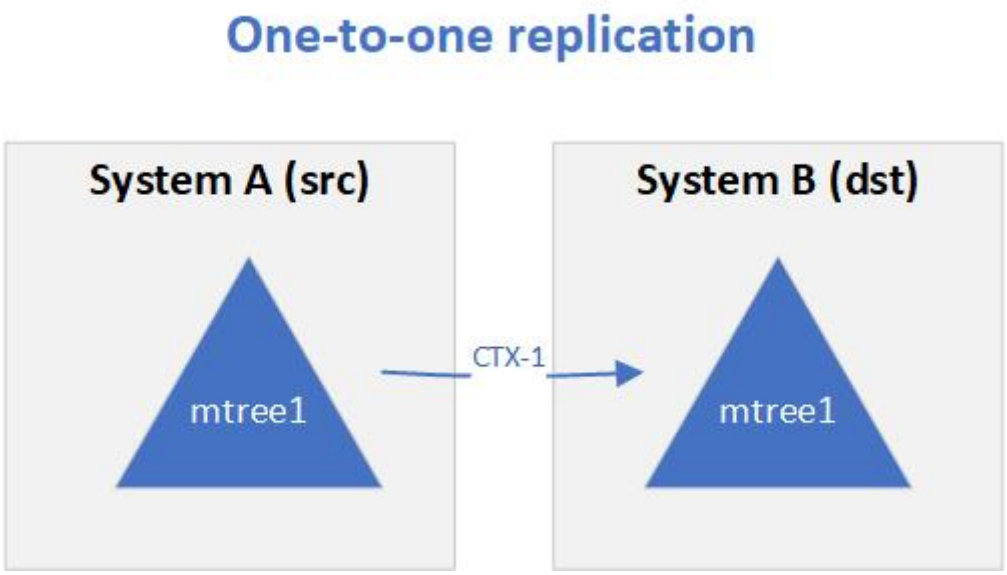

Figure 12. One-to-one replication pair

Related tasks

Creating a replication pair

Bi-directional replication

In a bi-directional replication pair, data from an MTree on DD system A is replicated to DD system B, and from another MTree on DD system B to DD system A.

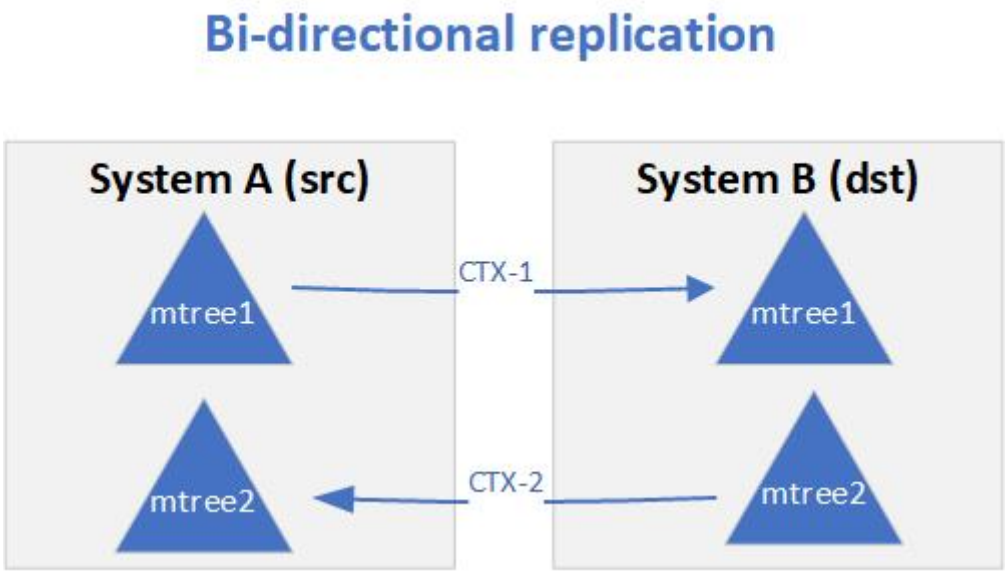

Figure 13. Bi-directional replication

Related concepts

Configuring bi-directional replication

One-to-many replication

In one-to-many replication, data flows from a source MTree on one DD system to several destination DD systems. You could use this type of replication to create more than two copies for increased data protection, or to distribute data for multi-site usage.

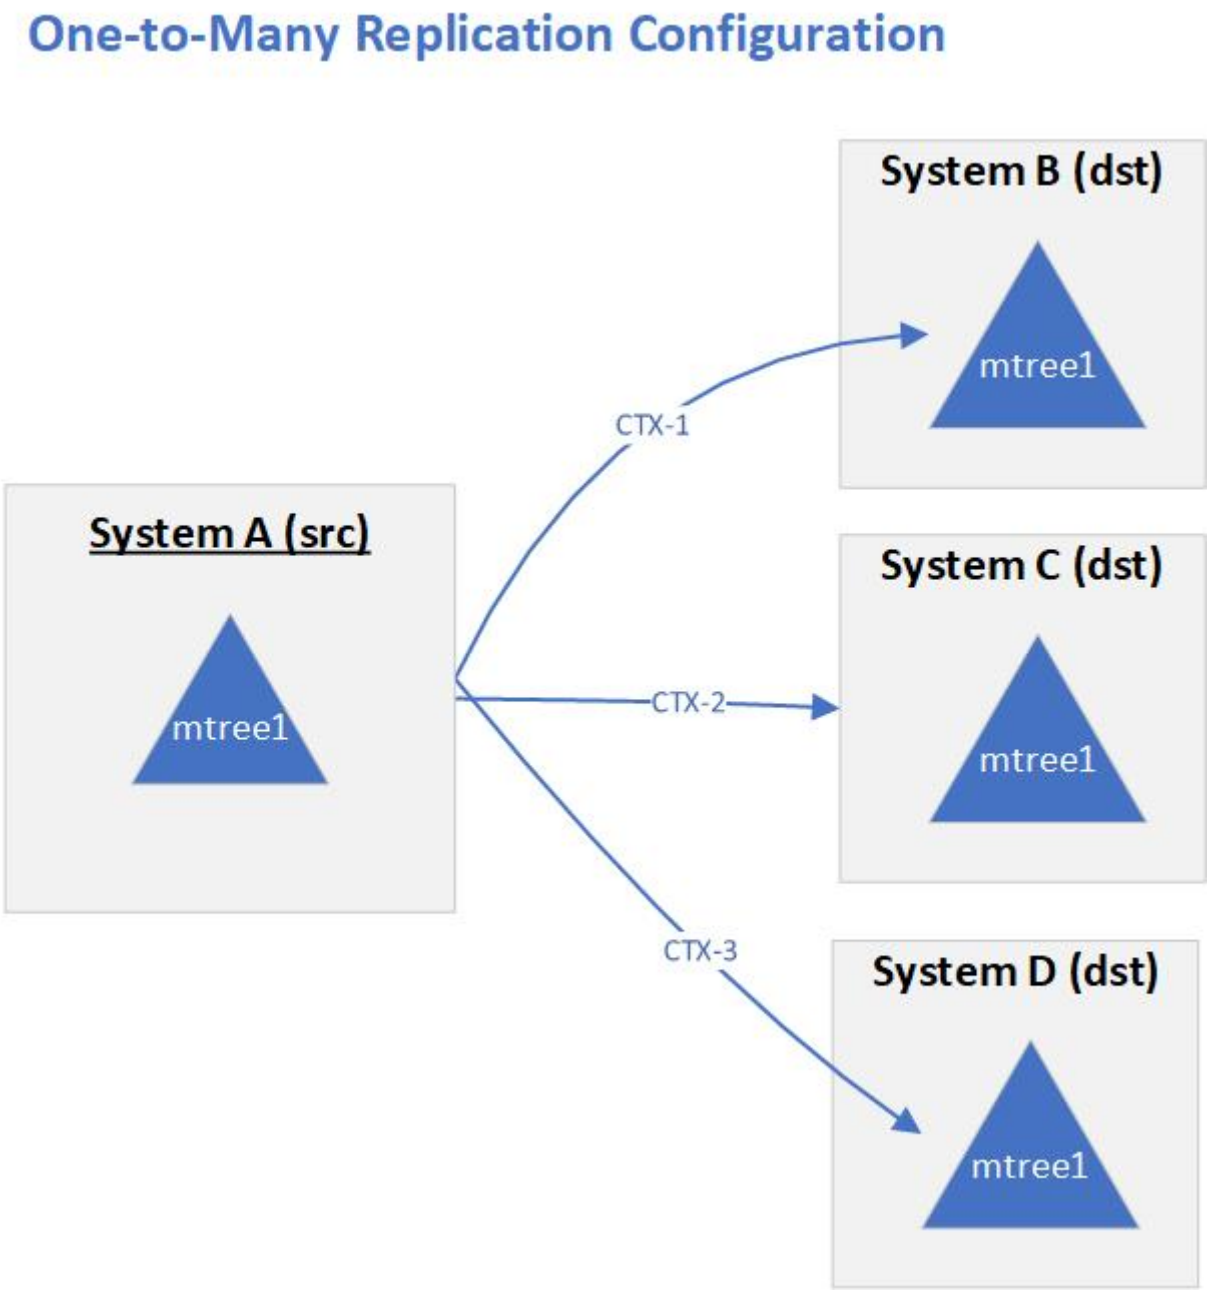

Figure 14. One-to-many replication

Related concepts

Configuring one-to-many replication

Many-to-one replication

In many-to-one replication, replication data flows from several source DD systems to a single destination DD system. This type of replication can be used to provide data recovery protection for several branch offices on a corporate headquarter's IT system.

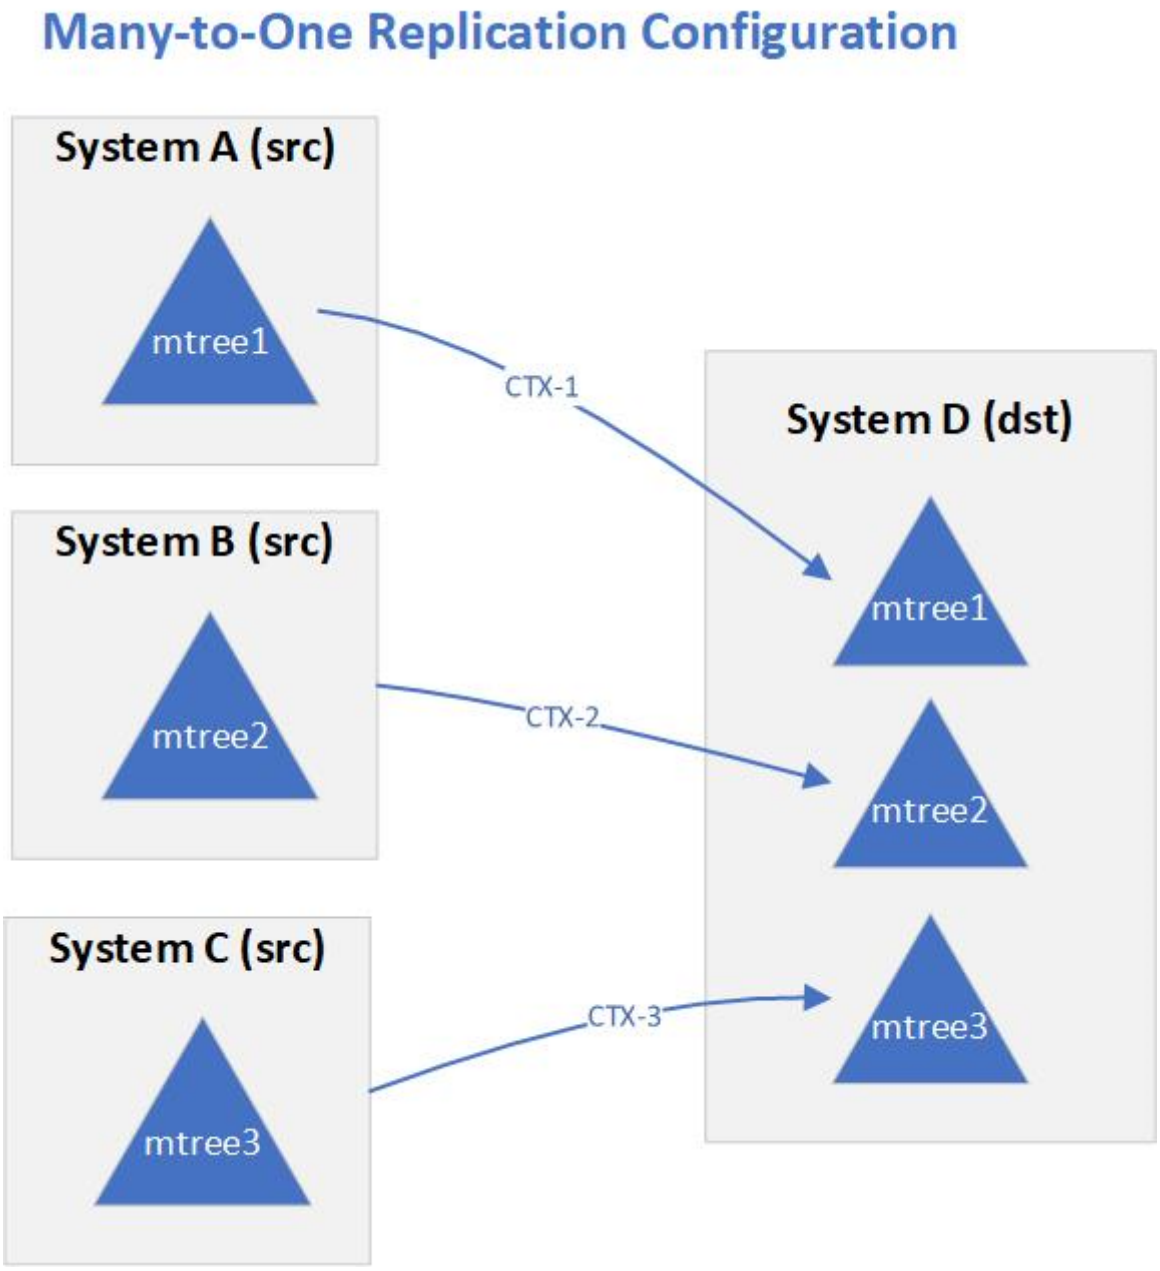

Figure 15. Many-to-one replication

Related concepts

Configuring many-to-one replication

Cascaded replication

In a cascaded replication topology, a source MTree is chained among three DD systems.

For example, DD system A replicates one or more MTrees to DD system B, which then replicates those MTrees to DD system C. The MTrees on DD system B are both a destination (from DD system A) and a source (to DD system C).

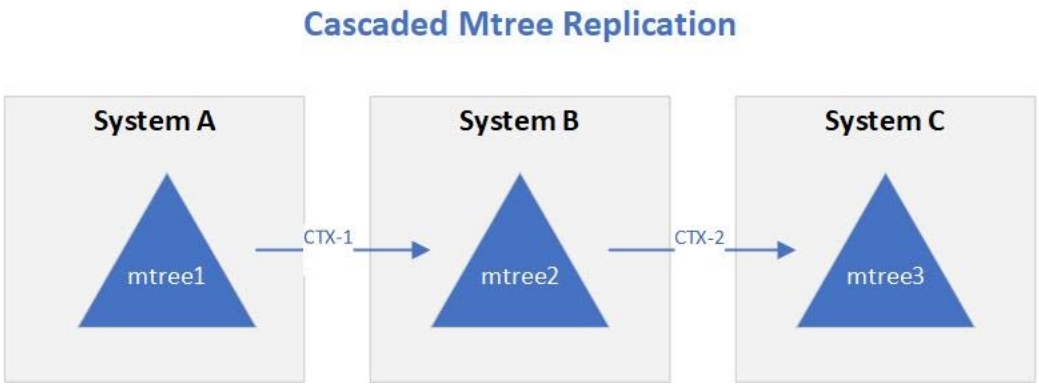

Figure 16. Cascade MTree replication

Data recovery for the corresponding MTree can be performed from the non-degraded replication pair context. For example:

- In the event a DD system A MTree requires recovery, data can be recovered from DD system B.

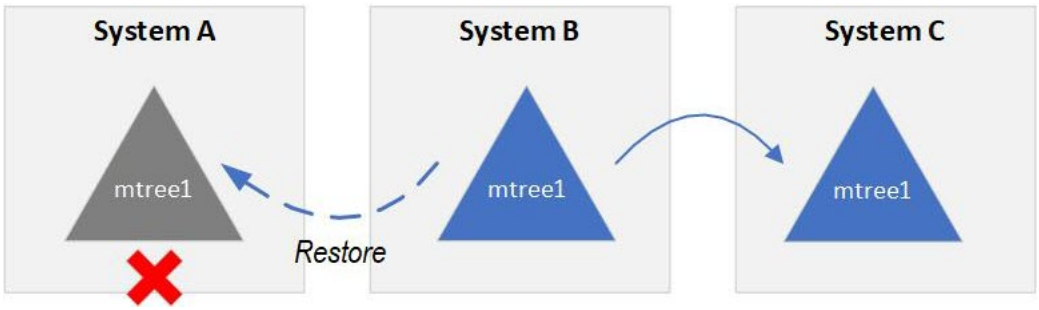

Figure 17. Cascade replication - Recovering MTree on system A

- In the event a DD system B MTree requires recovery, the simplest method is to perform an MTree replication resync from DD system A to DD system B. In this case, the replication context from DD system B to DD system C should be broken first. After the DD system A to DD system B replication context finishes resync, configure and resync a new DD system B to DD System C context.

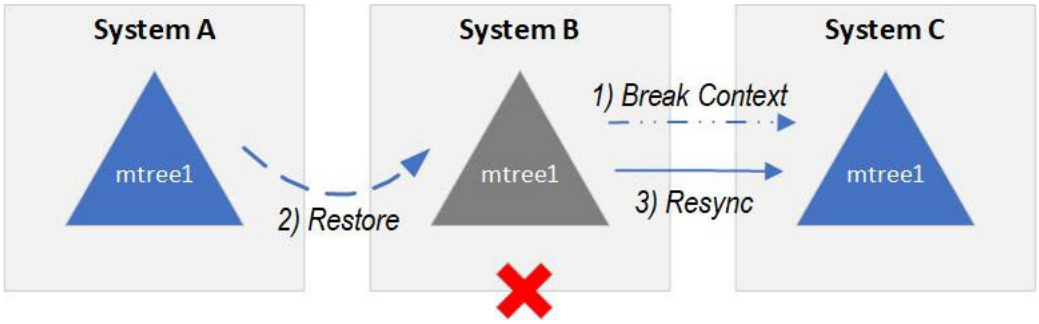

Figure 18. Cascade replication - Recovering MTree on system B

Related concepts

Configuring cascaded replication

Managing replication

You can manage replication using the DD System Manager) or the DDOS CLI.

About this task

To use a graphical user interface (GUI) to manage replication, log in to the DD System Manager.

Steps

1. From the menu at the left of the DD System Manager, select **Replication**. If your license has not been added yet, select **Add License**.
2. Select **Automatic** or **On-Demand** (you must have a DD Boost license for on-demand).

   CLI Equivalent

   You can also log in at the CLI:

```
login as: sysadmin
Data Domain OS 6.0.x.x-12345
Using keyboard-interactive authentication.
Password:
```

Replication status

Replication Status shows the system-wide count of replication contexts exhibiting a warning (yellow text) or error (red text) state, or if conditions are normal.

Summary view

The Summary view lists the configured replication contexts for a DD system, displaying aggregated information about the selected DD system – that is, summary information about the inbound and outbound replication pairs. The focus is the DD system, itself, and the inputs to it and outputs from it.

The Summary table can be filtered by entering a Source or Destination name, or by selecting a State (Error, Warning, or Normal).

Table 91. Replication Summary view

Item	Description
Source	System and path name of the source context, with format *system.path*. For example, for MTree `MTree1` on `system dd9900-22`, you would see `dd9900-22.chaos.local/data/col1/MTree1`.
Destination	System and path name of destination context, with format *system.path*. For example, for MTree `MTree1` on `system dd9900-44`, you would see `dd9900-44.chaos.local/data/col1/MTree1`.
Type	Type of context: MTree or Pool.
State	Possible states of replication pair status include: ● Normal – If the replica is Initializing, Replicating, Recovering, Resyncing, or Migrating. ● Idle – For MTree replication, this state can display if the replication process is not currently active or for network errors (such as the destination system being inaccessible). ● Warning – If there is an unusual delay for the first five states, or for the Uninitialized state. ● Error – Any possible error states, such as Disconnected.
Synced As Of Time	Timestamp for last automatic replication sync operation performed by the source. For MTree replication, this value is updated when a snapshot is exposed on the destination.
Pre-Comp Remaining	Amount of pre-compressed data remaining to be replicated.
Completion Time (Est.)	Value is either `Completed`, or the estimated amount of time required to complete the replication data transfer based on the last 24 hours' transfer rate.

Detailed information for a replication context

Selecting one replication context from the Summary view populates that context's information in Detailed Information, Performance Graph, Completion Stats, and Completion Predictor.

Table 92. Detailed Information

Item	Description
State Description	Message about state of replica.
Source	System and path name of source context, with format `system.path`. For example, for MTree `MTree1` on system `dd9900-22`, you would see `dd9900-22.chaos.local/data/col1/MTree1`.
Destination	System and path name of destination context, with format `system.path`. For example, for MTree `MTree1` on system `dd9900-44`, you would see `dd9900-44.chaos.local/data/col1/MTree1`.
Connection Port	System name and listen port used for replication connection.

Table 93. Performance Graph

Item	Description
Pre-Comp Remaining	Pre-compressed data remaining to be replicated.
Pre-Comp Written	Pre-compressed data written on the source.
Post-Comp Replicated	Post-compressed data that has been replicated.

Table 94. Completion Stats

Item	Description
Synced As Of Time	Timestamp for last automatic replication sync operation performed by the source. For MTree replication, this value is updated when a snapshot is exposed on the destination.
Completion Time (Est.)	Value is either `Completed` or the estimated amount of time required to complete the replication data transfer based on the last 24 hours' transfer rate.
Pre-Comp Remaining	Amount of data remaining to be replicated.
Status	For source and destination endpoints, shows status (Enabled, Disabled, Not Licensed, etc.) of major components on the system, such as: • Replication • File System • DD Retention Lock • DD Encryption at Rest • DD Encryption over Wire • Available Space • Low Bandwidth Optimization • Compression Ratio • Low Bandwidth Optimization Ratio

Completion Predictor

The Completion Predictor is a widget for tracking a backup job's progress and for predicting when replication will complete, for a selected context.

Related tasks

Changing host connection settings
Viewing estimated completion time for backup jobs

Creating a replication pair

Before creating a replication pair, make sure the destination does not *exist*, or you will get an error.

Steps

1. Select **Replication** > **Automatic** > **Summary tab** > **Create Pair** .
2. In the Create Pair dialog, add information to create an inbound or outbound MTree, collection, or pool replication pair, as described in the next sections.

Adding a DD system for replication

You may need to add a DD system as either a host or a destination before you can create a replication pair.

About this task

NOTE: Make sure the system being added is running a compatible DDOS version as described in Replication version compatibility.

Steps

1. In the Create Pair dialog, select Add System.
2. For System, enter the hostname or IP address of the system to be added.
3. For User Name and Password, enter the sysadmin's user name and password.
4. Optionally, select **More Options** to enter a proxy IP address (or system name) of a system that cannot be reached directly. If configured, enter a custom port instead of the default port 3009.

   NOTE: IPv6 addresses are supported only when adding a DDOS 5.5 or later system to a management system using DDOS 5.5 or later.

5. Select **OK**.

   NOTE: If the system is unreachable after adding it to DD System Manager, make sure that there is a route from the managing system to the system being added. If a hostname (either a fully qualified domain name (FQDN) or non-FQDN) is entered, make sure it is resolvable on the managed system. Configure a domain name for the managed system, ensure a DNS entry for the system exists, or ensure an IP address to hostname mapping is defined.

6. If the system certificate is not verified, the Verify Certificate dialog shows details about the certificate. Check the system credentials. Select **OK** if you trust the certificate, or select **Cancel**.

Creating a collection replication pair

See the *Collection replication* section for general information about this type of replication.

About this task

Before creating a collection replication pair, make sure:

- The storage capacity of the destination system is equal to, or greater than, that of the source system. (If the destination capacity is less than that of the source, the available capacity on the source is reduced to that of the destination.)
- The destination has been destroyed, and subsequently re-created, but not enabled.
- Each destination and each source is in only one context at a time.
- The file system is disabled on the replica, while configuring and enabling encryption on the source.
- The file system is disabled on the source, while configuring and enabling encryption on the replica.

Steps

1. In the Create Pair dialog, select **Collection** from the **Replication Type** menu.
2. Select the source system hostname from the **Source System** menu.
3. Select the destination system hostname from the **Destination System** menu. The list includes only those hosts in the DD-Network list.

4. If you want to change any host connection settings, select the **Advanced** tab.
5. Select **OK**. Replication from the source to the destination begins.

Results

Test results returned the following performance guidelines for replication initialization. These are guidelines *only*, and actual performance seen in production environments may vary.

- Over a gibibit LAN: With a high enough shelf count to drive maximum input/output and ideal conditions, collection replication can saturate a 1GigE link (modulo 10% protocol overhead), as well as 400-900 MB/sec on 10gigE, depending on the platform.
- Over a WAN, performance is governed by the WAN link line speed, bandwidth, latency, and packet loss rate.

Related tasks

Changing host connection settings

Creating an MTree or pool replication pair

See the *MTree replication* section for general information about this type of replication.

About this task

When creating an MTree or pool replication pair:

- Make sure the replication is transiting\exiting the correct interface. When defining a replication context, the host names of the source and destination must resolve with forward and reverse lookups. To make the data transit alternate interfaces on the system, other than the default resolving interface, the replication context must be modified after creation. It may be necessary to set up host files to ensure that contexts are defined on non-resolving (cross-over) interfaces.
- You can "reverse" the context for an MTree replication, that is, you can switch the destination and the source.
- Subdirectories within an MTree cannot be replicated, because the MTree, in its entirety, is replicated.
- The destination DD system must have available storage capacity of at least the post-compressed size of the expected maximum post-compressed size of the source directory or MTree.
- When replication is initialized, a destination directory is created automatically.
- A DD system can simultaneously be the source for one context and the destination for another context.

Steps

1. In the Create Pair dialog, select **MTree** (default), or **Pool** from the **Replication Type** menu.
2. Select the source system hostname from the **Source System** menu.
3. Select the destination system hostname from the **Destination System** menu.
4. Enter the source path in the **Source Path** text box (notice the first part of the path is a constant that changes based on the type of replication chosen).
5. Enter the destination path in the **Destination Path** text box (notice the first part of the path is a constant that changes based on the type of replication chosen).
6. If you want to change any host connection settings, select the **Advanced** tab.
7. Select **OK**.

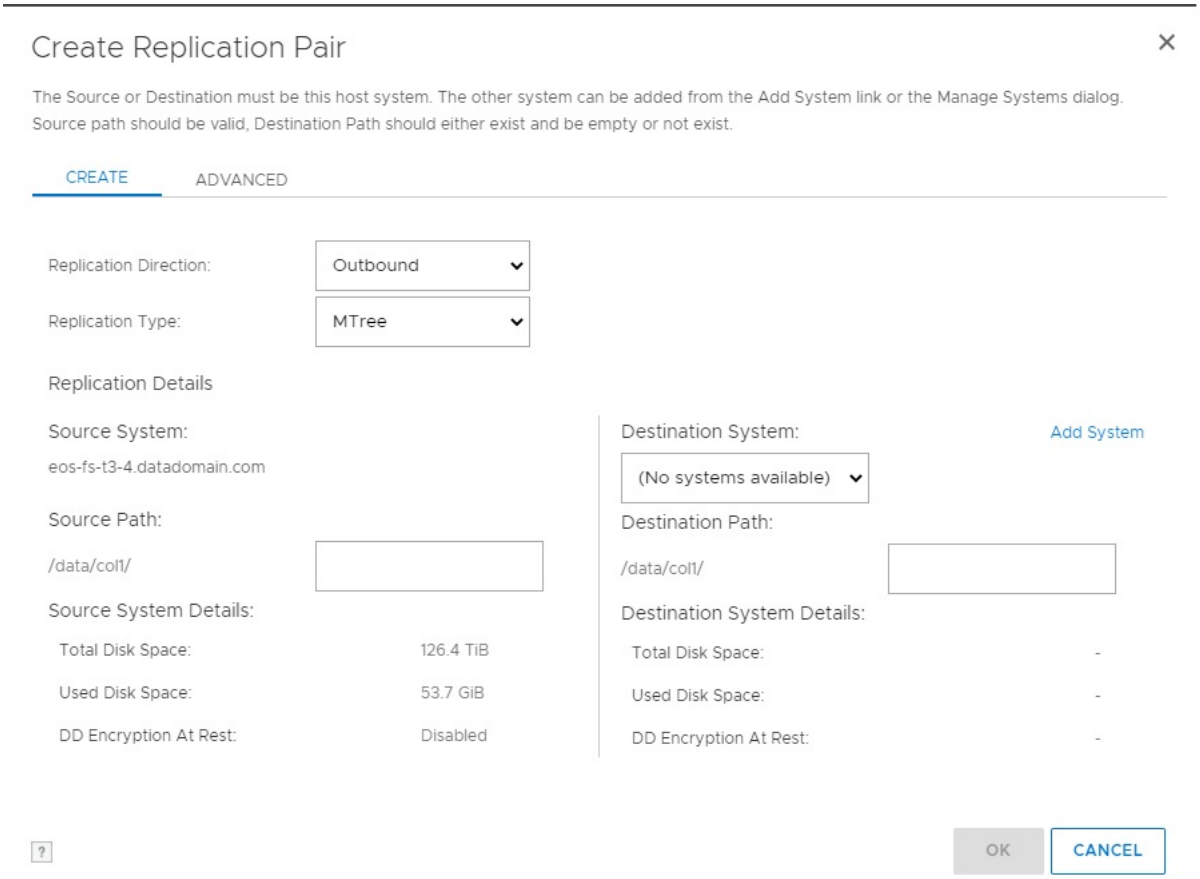

Figure 19. Creating an MTree or pool replication pair

The Replication from the source to the destination begins.

Test results from returned the following guidelines for estimating the time needed for replication initialization.

These are guidelines *only* and may not be accurate in specific production environments.

- Using a T3 connection, 100ms WAN, performance is about 40 MiB/sec of pre-compressed data, which gives data transfer of:

  40 MiB/sec = 25 seconds/GiB = 3.456 TiB/day

- Using the base-2 equivalent of gigabit LAN, performance is about 80 MiB/sec of pre-compressed data, which gives data transfer of about double the rate for a T3 WAN.

CLI Equivalent

Here is an example of creating MTree replication pairs at the CLI. In this example, the source system is **`dd-src`** and the destination system is **`dd-dst`**. For details about usage in other scenarios, see the *DDOS Command Reference Guide*.

1. Create an MTree on the source system:

```
sysadmin@dd-src mtree create /data/col1/mtree-src
MTree "/data/col1/mtree-src" created successfully.
```

2. Create the replication context in the source, using the full hostname.

```
sysadmin@dd-dst replication add source mtree://dd-src.domain.com/data/col1/mtree-src
destination mtree://dd-dst.domain.com/data/col1/mtree-dst
```

3. Create the replication context in the destination system, using the full hostname.

```
sysadmin@dd-dst replication add source mtree://dd-src.domain.com/data/col1/mtree-src
destination mtree://dd-dst.domain.com/data/col1/mtree-dst
```

4. To verify that the MTree replication context has been created, use the `replication show config` command.

   The output is horizontally truncated in this example.

```
sysadmin@dlh5# replication show config
CTX  Source                                          Destination
---  ----------------------------------------
-----------------------------------------
1    mtree://dd-src.domain.com/data/col1/mtree-src mtree://dd-dst.domain.com/data/col1/
mtree-dst
---  ----------------------------------------
-----------------------------------------
 * Used for recovery only.
```

5. To start replication between a source and destination, use the `replication initialize` command on the source. This command checks that the configuration and connections are correct and returns error messages if any problems occur.

```
sysadmin@dd-src# replication initialize mtree://dd-dst.domain.com/data/col1/mtree-dst
 (00:08) Waiting for initialize to start...
 (00:10) Intialize started.
 Use 'replication watch mtree://dd-dst.domain.com/data/col1/mtree-dst' to monitor
progress.
```

Related tasks

Changing host connection settings

Configuring bi-directional replication

To create a bi-directional replication pair, use the MTree replication pair procedure (for example, using mtree2) from host A to host B. Use the same procedure to create a replication pair (for example, using mtree1) from host B to host A. For this configuration, destination pathnames cannot be the same.

Related tasks

Creating an MTree or pool replication pair

Configuring one-to-many replication

To create a one-to-many replication pair, use the MTree replication pair procedure (for example, using mtree1) on host A to: (1) mtree1 on host B, (2) mtree1 on host C, and (3) mtree1 on host D. A replication recovery cannot be done to a source context whose path is the source path for other contexts; the other contexts must be broken and resynced after the recovery.

Related tasks

Creating an MTree or pool replication pair

Configuring many-to-one replication

To create a many-to-one replication pair, use the MTree replication pair procedure [for example, (1) mtree1 from host A to mtree1 on host C and (2) mtree2 on host B to mtree2 on host C.]

Related tasks

Creating an MTree or pool replication pair

Configuring cascaded replication

To create a cascaded replication pair, use the MTree replication pair procedure: (1) mtree1 on host A to mtree1 on host B, and (2) on host B, create a pair for mtree1 to mtree1 on host C. The final destination context (on host C in this example, but more than three hops are supported) can be a collection replica or MTree replica.

Related tasks

Creating an MTree or pool replication pair

Creating a schedule for an MTree replication pair

For MTree replication contexts, DDOS supports the creation of a schedule to automatically enable and disable the replication context. This feature is only available through CLI commands.

Steps

1. Run the **replication schedule set *context* enable *hhmm* disable *hhmm*** command to set the enable and disable times for the specified MTree replication context.
2. To change the schedule later, run the **replication schedule set *context* enable *hhmm* disable *hhmm*** command again for the same context and specify the new enable and disable times.
3. To delete a schedule , run the **replication schedule reset *context* enable *hhmm* disable *hhmm*** command and specify the context, enable time, and disable time to reset.
4. To display a list of replication schedules, run the **replication schedule show** command.

Disabling and enabling a replication pair

Disabling a replication pair temporarily pauses the active replication of data between a source and a destination. The source stops sending data to the destination, and the destination stops serving as an active connection to the source.

Steps

1. Select one or more replication pairs in the Summary table, and select **Disable Pair**.
2. In the Display Pair dialog, select **Next** and then **OK**.
3. To resume operation of a disabled replication pair, select one or more replication pairs in the Summary table, and select **Enable Pair** to display the Enable Pair dialog.
4. Select **Next** and then **OK**. Replication of data is resumed.

CLI Equivalent

```
# replication disable {destination | all}

# replication enable {destination | all}
```

Deleting a replication pair

When an MTree replication pair is deleted, the destination MTree becomes writable. When a collection replication pair is deleted, the destination DD system becomes a stand-alone read/write system, and the file system is disabled.

Steps

1. Select one or more replication pairs in the Summary table, and select **Delete Pair**.
2. In the Delete Pair dialog, select **Next** and then **OK**. The replication pairs are deleted.

CLI Equivalent

Before running this command, always run the `filesys disable` command. Then, afterward, run the `filesys enable` command

```
# replication break {destination | all}
```

Changing host connection settings

To direct traffic out of a specific port, modify a current context by altering the connection host parameter using a host name previously defined in the local hosts file to address the alternate system. That host name will correspond to the destination. The host entry will indicate an alternate destination address for that host. This may be required on both the source and destination systems.

Steps

1. Select the replication pair in the Summary table, and select **Modify Settings**. You can also change these settings when you are performing Create Pair, Start Resync, or Start Recover by selecting the **Advanced** tab.
2. In the Modify Connection Settings dialog, modify any or all of these settings:
   a. **Use Low Bandwidth Optimization** – For enterprises with small data sets and 6 Mb/s or less bandwidth networks, DD Replicator can further reduce the amount of data to be sent using *low bandwidth optimization*. This enables remote sites with limited bandwidth to use less bandwidth or to replicate and protect more of their data over existing networks. Low bandwidth optimization must be enabled on both the source and destination DD systems. If the source and destination have incompatible low bandwidth optimization settings, low bandwidth optimization will be inactive for that context. After enabling low bandwidth optimization on the source and destination, both systems must undergo a full cleaning cycle to prepare the existing data, so run `filesys clean start` on both systems. The duration of the cleaning cycle depends on the amount of data on the DD system, but takes longer than a normal cleaning. For more information on the `filesys` commands, see the *DDOS Command Reference Guide*.

      Important: Low bandwidth optimization is not supported for Collection Replication.
   b. **Enable Encryption Over Wire** – DD Replicator supports encryption of data-in-flight with TLS protocol version 1.1. When replication authentication-mode is set to **`one-way`** or **`two-way`**, DHE (Ephemeral Diffie-Hellman) is used for session key exchange. Server authentication happens via RSA. The AES 256-bit GCM cipher encapsulates the replicated data over the wire. The encryption encapsulation layer is immediately removed as soon as it lands on the destination system. SHA384 is used for hash/Message Authentication Code.

      One-way indicates that only the destination certificate is verified. Two-way indicates that both the source and destination certificates are verified. Mutual trust must be established before you can use the `authentication-mode` option, and both sides of the connection must enable this feature for encryption to proceed.

      When the replication authentication mode is set to **`anonymous`**, ADH (Anonymous Diffie-Hellman) is used for session key exchanges, but the source and destination do not authenticate each other before the key exchange . If the authentication mode is not specified, anonymous is the default value.
   c. **Network Preference** – You may choose IPv4 or IPv6. An IPv6-enabled replication service can still accept connections from an IPv4 replication client if the service is reachable via IPv4. An IPv6-enabled replication client can still communicate with an IPv4 replication service if the service is reachable via IPv4.
   d. **Use Non-default Connection Host** – The source system transmits data to a destination system listen port. Since a source system can have replication configured for many destination systems (each of which can have a different listen port), each context on the source can configure the connection port to the corresponding listen port of the destination.
3. Select **Next** and then **Close**.

   The replication pair settings are updated, and replication resumes.

   CLI Equivalent

```
# replication modify <destination> connection-host <new-host-name> [port <port>]
#  replication modify <destination> connection-host <new-host-name> [port <port>]
#  replication modify <destination> low-bw-optim {enabled | disabled}
#  replication modify <destination> encryption {enabled [authentication-mode {one-way |
two-way | anonymous}] | disabled}
#  replication modify <destination> ipversion {ipv4 | ipv6}
```

Managing replication systems

You can add or delete protection systems to be used for replication using the Manage Systems dialog.

Steps

1. Select **Manage Systems**.
2. In the Manage Systems dialog, add and/or delete systems, as required.
3. Select **Close**.

Recovering data from a replication pair

If source replication data becomes inaccessible, it can be *recovered* from the replication pair destination. The source must be empty before recovery can proceed. Recovery can be performed for all replication topologies, except for MTree replication.

Recovery of data from a directory pool, as well as from collection replication pairs, is described in the next sections.

Recovering directory pool data

You can recover data from a directory-based pool, but not from an MTree-based pool.

Steps

1. Select **More** > **Start Recover**.
2. In the Start Recover dialog, select **Pool** from the **Replication Type** menu.
3. Select the source system hostname from the **System to recover to** menu.
4. Select the destination system hostname from the **System to recover from** menu.
5. Select the context on the destination from which data is recovered.
6. If you want to change any host connection settings, select the **Advanced** tab.
7. Select **OK** to start the recovery.

Related tasks

Changing host connection settings

Recovering collection replication pair data

To successfully recover collection replication pair data, the source file system must be in a pristine state, and the destination context must be fully initialized.

Steps

1. Select **More** > **Start Recover** to display the Start Recover dialog.
2. Select **Collection** from the **Replication Type** menu.
3. Select the source system host name from the **System to recover to** menu.
4. Select the destination system host name from the **System to recover from** menu.
5. Select the context on the destination from which data is recovered. Only one collection will exist on the destination.
6. To change any host connection settings, select the **Advanced** tab.
7. Select **OK** to start the recovery.

Related tasks

Changing host connection settings

Aborting a replication pair recovery

If a replication pair recovery fails or must be terminated, you can stop the replication recovery.

Steps

1. Select the More menu, and select **Abort Recover** to display the Abort Recover dialog, which shows the contexts currently performing recovery.
2. Select the checkbox of one or more contexts to abort from the list.
3. Select **OK**.

Next steps

As soon as possible, you should restart recovery on the source.

Resyncing an MTree or pool replication pair

Resynchronization is the process of recovering (or bringing back into sync) the data between a source and a destination replication pair after a manual break. The replication pair are resynchronized so both endpoints contain the same data. Resynchronization is available for MTree and pool replication, but not for collection replication.

About this task

A replication resynchronization can also be used:

- To recreate a context that has been deleted.
- When a destination runs out of space, but the source still has data to replicate.

Steps

1. Delete the context on both the replication source and replication destination systems.
2. From either the replication source or replication destination system, select **More** > **Start Resync** to display the Start Resync dialog.
3. Select the Replication Type to be resynced: **MTree**, or **Pool**.
4. Select the replication source system host name from the **Source System** menu.
5. Select the replication destination system host name from the **Destination System** menu.
6. Enter the replication source path in the **Source Path** text box.
7. Enter the replication destination path in the **Destination Path** text box.
8. To change any host connection settings, select the **Advanced** tab.
9. Select **OK**.

CLI Equivalent

```
# replication resync destination
```

Related tasks

Changing host connection settings

Aborting a replication pair resynchronization

If a replication pair resynchronization fails or must be terminated, you can stop the resynchronization.

Steps

1. From either the replication source or replication destination system, select **More** > **Abort Resync** to display the Abort Resync dialog, which lists all contexts currently performing resynchronization.
2. Select the checkboxes of one or more contexts to abort their resynchronization.
3. Select **OK**.

DD Boost view

The DD Boost view provides configuration and troubleshooting information to NetBackup administrators who have configured DD systems to use DD Boost AIR (Automatic Image Replication) or any DD Boost application that uses managed file replication.

See the *DD Boost for OpenStorage Administration Guide* for DD Boost AIR configuration instructions.

The **File Replication** tab displays:

- Currently Active File Replication:
  - Direction (Out-Going and In-Coming) and the number of files in each.
  - Remaining data to be replicated (pre-compressed value in GiB) and the amount of data already replicated (pre-compressed value in GiB).
  - Total size: The amount of data to be replicated and the already replicated data (pre-compressed value in GiB).
- Most Recent Status: Total file replications and whether completed or failed

    - during the last hour
    - over the last 24 hours
- Remote Systems:
    - Select a replication from the list.
    - Select the time period to be covered from the menu.
    - Select **Show Details** for more information about these remote system files.

The **Storage Unit Associations** tab displays the following information, which you can use for audit purposes or to check the status of DD Boost AIR events used for the storage unit's image replications:

- A list of all storage unit **Associations** known to the system. The source is on the left, and the destination is on the right. This information shows the configuration of AIR on the protection system.
- The **Event Queue** is the pending event list. It shows the local storage unit, the event ID, and the status of the event.

An attempt is made to match both ends of a DD Boost path to form a pair and present this as one pair/record. If the match is impossible, for various reasons, the remote path will be listed as *Unresolved*.

Remote system files

The Show Details button provides information for the selected remote file replication system. File Replications shows starting and ending information, as well as size and data amount, for the selected remote file replication system. The Performance Graph shows performance over time for the selected remote file replication system.

Table 95. File Replications

Item	Description
Start	Starting point of time period.
End	Ending point of time period.
File Name	Name of specific replication file.
Status	Most recent status (Success, Failure).
Pre-Comp Size (MiB)	Amount of pre-compressed outbound and inbound data, as compared to network throughput or post-compressed data (in MiB).
Network Bytes (MiB)	Amount of network throughput data (in MiB).

Table 96. Performance Graph

Item	Description
Duration	Duration for replication (either 1d, 7d or 30d).
Interval	Interval for replication (either Daily or Weekly).
Pre-Comp Replicated	Amount of pre-compressed outbound and inbound data (in GiB).
Post-Comp Replicated	Amount of post-compressed data (in GiB).
Network Bytes	Amount of network throughput data (in GiB).
Files Succeeded	Number of files that were successfully replicated.
Files Failed	Number of files that failed to be replicated.
Show in new window	Brings up a separate window.
Print	Prints the graph.

Performance view

The Performance view displays a graph that represents the fluctuation of data during replication. These are aggregated statistics of each replication pair for this DD system.

- **Duration** (x-axis) is 30 days by default.
- **Replication Performance** (y-axis) is in GibiBytes or MebiBytes (the binary equivalents of GigaBytes and MegaBytes).
- **Network In** is the total replication network bytes entering the system (all contexts).

- **Network Out** is the total replication network bytes leaving the system (all contexts).
- For a reading of a specific point in time, hover the cursor over a place on the graph.
- During times of inactivity (when no data is being transferred), the shape of the graph may display a gradually descending line, instead of an expected sharply descending line.

Advanced Settings view

Advanced Settings lets you manage throttle and network settings.

Throttle Settings

- **Throttle Override** – Displays throttle rate if configured, or 0 meaning all replication traffic is stopped.
- **Permanent Schedule** – Displays the time and days of the week on which scheduled throttling occurs.

Network Settings

- **Bandwidth** – Displays the configured data stream rate if bandwidth has been configured, or Unlimited (default) if not. The average data stream to the replication destination is at least 98,304 bits per second (12 KiB).
- **Delay** – Displays the configured network delay setting (in milliseconds) if it has been configured, or None (default) if not.
- **Listen Port** – Displays the configured listen port value if it has been configured, or 2051 (default) if not.

Adding throttle settings

To modify the amount of bandwidth used by a network for replication, you can set a *replication throttle* for replication traffic.

About this task

There are three types of replication throttle settings:

- **Scheduled throttle** – The throttle rate is set at a predetermined time or period.
- **Current throttle** – The throttle rate is set until the next scheduled change, or until a system reboot.
- **Override throttle** – The previous two types of throttle are overridden. This persists – even through reboot – until you select **Clear Throttle Override** or issue the `replication throttle reset override` command.

You can also set a default throttle or a throttle for specific destinations, as follows:

- **Default throttle** – When configured, all replication contexts are limited to this throttle, except for those destinations specified by destination throttles (see next item).
- **Destination throttle** – This throttle is used when only a few destinations need to be throttled, or when a destination requires a throttle setting different from the default throttle. When a default throttle already exists, this throttle takes precedence for the destination specified. For example, you can set the default replication throttle to *10 kbps*, but – using a destination throttle – you can set a single collection replication context to *unlimited*.

  NOTE: Currently, you can set and modify destination throttle only by using the command-line interface (CLI); this functionality is not available in the DD System Manager. For documentation on this feature, see the `replication throttle` command in the *DDOS Command Reference Guide*. If the DD System Manager detects that you have one or more destination throttles set, you will be given a warning, and you should use the CLI to continue.

Additional notes about replication throttling:

- Throttles are set only at the source. The only throttle that applies to a destination is the **0 Bps (Disabled)** option, which disables all replication traffic.
- The minimum value for a replication throttle is 98,304 bits per second.

Steps

1. Select **Replication** > **Advanced Settings** > **Add Throttle Setting** to display the Add Throttle Setting dialog.
2. Set the days of the week for which throttling is to be active by selecting **Every Day** or by selecting checkbox(es) next to individual day(s).
3. Set the time that throttling is to start with the **Start Time** drop-down selectors for the hour:minute and AM/PM.
4. For **Throttle Rate**:
   - Select **Unlimited** to set no limits.
   - Enter a number in the text box (for example, 20000), and select the rate from the menu (bps, Kbps, Bps, or KBps).

- Select the **0 Bps (disabled)** option to disable all replication traffic.

5. Select **OK** to set the schedule. The new schedule is shown under **Permanent Schedule**.

Results

Replication runs at the given rate until the next scheduled change, or until a new throttle setting forces a change.

Deleting Throttle Settings

You can delete a single throttle setting or all throttle settings at once.

Steps

1. Select **Replication** > **Advanced Settings** > **Delete Throttle Setting** to display the Delete Throttle Setting dialog.
2. Select the checkbox for the throttle setting to delete, or select the heading checkbox to delete all settings. This list can include settings for the "disabled" state.
3. Select **OK** to remove the setting.
4. In the Delete Throttle Setting Status dialog, select **Close**.

Temporarily overriding a throttle setting

A throttle override temporarily changes a throttle setting. The current setting is listed at the top of the window.

Steps

1. Select **Replication** > **Advanced Settings** > **Set Throttle Override** to display the Throttle Override dialog.
2. Either set a new throttle override, or clear a previous override.
   a. To set a new throttle override:
      - Select **Unlimited** to revert to the system-set throttle rate (no throttling performed), or
      - Set the throttling bit and rate in the text box (for example, 20000) and (bps, Kbps, Bps, or KBps), or
      - Select **0 Bps (Disabled)** to set the throttle rate to 0, effectively stopping all replication network traffic.
      - To enforce the change temporarily, select **Clear at next scheduled throttle event**.

   b. To clear an override previously set, select **Clear Throttle Override**.
3. Select **OK**.

Changing network settings

Using the bandwidth and network-delay settings together, replication calculates the proper TCP (transmission control protocol) buffer size for replication usage. These network settings are global to the DD system and should be set only once per system.

About this task

Note the following:

- You can determine the actual bandwidth and the actual network delay values for each server by using the `ping` command.
- The default network parameters in a restorer work well for replication in low latency configurations, such as a local 100Mbps or 1000Mbps Ethernet network, where the latency round-trip time (as measured by the `ping` command) is usually less than 1 millisecond. The defaults also work well for replication over low- to moderate-bandwidth WANs, where the latency may be as high as 50-100 milliseconds. However, for high-bandwidth high-latency networks, some tuning of the network parameters is necessary.

  The key number for tuning is the bandwidth-delay number produced by multiplying the bandwidth and round-trip latency of the network. This number is a measure of how much data can be transmitted over the network before any acknowledgments can return from the far end. If the bandwidth-delay number of a replication network is more than 100,000, then replication performance benefits from setting the network parameters in both restorers.

Steps

1. Select **Replication** > **Advanced Settings** > **Change Network Settings** to display the Network Settings dialog.

2. In the Network Settings area, select **Custom Values**.
3. Enter **Delay** and **Bandwidth** values in the text boxes. The network delay setting is in milliseconds, and bandwidth is in bytes per second.
4. In the Listen Port area, enter a new value in the text box. The default IP Listen Port for a replication destination for receiving data streams from the replication source is 2051. This is a global setting for the DD system.
5. Select **OK**. The new settings appear in the Network Settings table.

Monitoring replication

The DD System Manager provides many ways to track the status of replication – from checking replication pair status, to tracking backup jobs, to checking performance, to tracking a replication process.

Viewing estimated completion time for backup jobs

You can use the Completion Predictor to see the estimated time for when a backup replication job will be completed.

Steps

1. Select **Replication** > **Summary**.
2. Select a Replication context for which to display Detailed Information.
3. In the Completion Predictor area, select options from the **Source Time** drop-down list for a replication's completion time, and select **Track**.

   The estimated time displays, in the Completion Time area, for when a particular backup job will finish its replication to the destination. If the replication is finished, the area shows `Completed`.

Checking replication context performance

To check the performance of a replication context over time, select a Replication context in the Summary view, and select **Performance Graph** in the Detailed Information area.

Tracking status of a replication process

To display the progress of a replication initialization, resynchronization, or recovery operation, use the **Replication** > **Summary** view to check the current state.

CLI Equivalent

When specifying an IP version, use the following command to check its setting:

```
sysadmin@dd-src# replication show config all

CTX:                          1
Source:                       mtree://dd-src.domain.com/data/col1/mtree-src
Destination:                  mtree://dd-dst.domain.com/data/col1/mtree-dst
Connection Host:              dd-src.domain.com*
Connection Port:              (default)*
Ipversion:                    (default)
Low-bw-optim:                 disabled
Crepl-gc-bw-optim             disabled
Encryption:                   disabled
Max-repl-streams:             -
Enabled:                      yes
Propagate-retention-lock:     disabled
```

Replication lag

The amount of time between two copies of data is known as replication lag.

You can measure the replication lag between two contexts with the replication status command. For information about determining the cause of replication lag and mitigating its impact, see the KB article *Data Domain: Troubleshooting Replication Lag*, available from the Online Support website.

Replication with HA

Floating IP addresses allow HA systems to specify a single IP address for replication configuration that will work regardless of which node of the HA pair is active.

Over IP networks, HA systems use a floating IP address to provide data access to the HA pair, regardless of which physical node is the active node. The net config command provides the **[type {fixed | floating}]** option to configure a floating IP address. The *DDOS Command Reference Guide* provides more information.

If a domain name is needed to access the floating IP address, specify the HA system name as the domain name. Run the `ha status` command to locate the HA system name.

NOTE: Run the `net show hostname type ha-system` command to display the HA system name, and if required, run the `net set hostname ha-system command` to change the HA system name.

All file system access should be through the floating IP address. When configuring backup and replication operations on an HA pair, always specify the floating IP address as the IP address for the protection system. Other system features such as DD Boost and replication will accept the floating IP address for the HA pair the same way as they accept the system IP address for a non-HA system.

Replication between HA and non-HA systems

Collection replication between HA and non-HA systems is not supported. MTree replication is required to replicate data between HA and non-HA systems.

Replicating a system with quotas to one without

Replicate a system with a DDOS that supports quotas, to a system with a DDOS that does not have quotas.

- A reverse resync, which takes the data from the system without quotas and puts it back in an MTree on the system that has quotas enabled (and which continues to have quotas enabled).
- A reverse initialization from the system without quotas, which takes its data and creates a new MTree on the system that supports quotas, but does not have quotas enabled because it was created from data on a system without quotas.

Replication Scaling Context

The Replication Scaling Context feature gives you more flexibility when configuring replication contexts.

In environments with more than 299 MTree replication contexts, this feature allows you to configure the contexts in any order.

The total number of replication contexts cannot exceed 540.

Using collection replication for disaster recovery with SMT

To use the destination system of a collection replication pair configured with SMT as a replacement system for disaster recovery, additional SMT configuration steps must be performed in addition to the other configuration steps required to bring a replacement system online.

Prerequisites

Using the collection replication destination system in this manner requires autosupport reports to be configured and saved. The KB article *Collection replica with smt enabled*, available from the Online Support website, provides additional information.

About this task

The replacement system will not have the following SMT details:

- Alert notification lists for each tenant-unit
- All users assigned to the DD Boost protocol for use by SMT tenants, if DD Boost is configured on the system
- The default-tenant-unit associated with each DD Boost user, if any, if DD Boost is configured on the system

Complete the following steps to configure SMT on the replacement system.

Steps

1. In the autosupport report, locate the output for the `smt tenant-unit show detailed` command.

```
Tenant-unit: "tu1"
Summary:
Name    Self-Service    Number of Mtrees    Types           Pre-Comp(GiB)
----    ------------    ----------------    --------        -------------
tu1     Enabled         2                   DD Boost        2.0
----    ------------    ----------------    --------        -------------

Management-User:
User      Role
------    ------------
tu1_ta    tenant-admin
tu1_tu    tenant-user
tum_ta    tenant-admin
------    ------------

Management-Group:
Group     Role
------    ------------
qatest    tenant-admin
------    ------------

DDBoost:
Name    Pre-Comp (GiB)    Status    User     Tenant-Unit
----    --------------    ------    -----    -----------
su1               2.0     RW/Q      ddbu1    tu1
----    --------------    ------    -----    -----------
Q    : Quota Defined
RO   : Read Only
RW   : Read Write

Getting users with default-tenant-unit tu1
DD Boost user    Default tenant-unit
-------------    -------------------
ddbu1            tu1
-------------    -------------------

Mtrees:
Name              Pre-Comp (GiB)    Status    Tenant-Unit
--------------    --------------    ------    -----------
/data/col1/m1                0.0    RW/Q      tu1
/data/col1/su1               2.0    RW/Q      tu1
--------------    --------------    ------    -----------
D : Deleted
```

```
Q : Quota Defined
RO : Read Only
RW : Read Write
RD : Replication Destination
IRH : Retention-Lock Indefinite Retention Hold Enabled
ARL : Automatic-Retention-Lock Enabled
RLGE : Retention-Lock Governance Enabled
RLGD : Retention-Lock Governance Disabled
RLCE : Retention-Lock Compliance Enabled


Quota:
Tenant-unit: tu1
Mtree            Pre-Comp (MiB)   Soft-Limit (MiB)   Hard-Limit(MiB)
--------------   --------------   ----------------   ----------------
/data/col1/m1                 0              71680             81920
/data/col1/su1             2048              30720             51200
--------------   --------------   ----------------   ----------------

Alerts:
Tenant-unit: "tu1"
Notification list "tu1_grp"
Members
------------------
tom.tenant@abc.com
------------------

No such active alerts.
```

2. On the replacement system, enable SMT if it is not already enabled.
3. On the replacement system, license and enable DD Boost if it is required and not already enabled.
4. If DD Boost is configured, assign each user listed in the `DD Boost` section of the "smt tenant-unit show detailed" output as a DD Boost User.

```
# ddboost user assign ddbu1
```

5. If DD Boost is configured, assign each user listed in the `DD Boost` section of the `smt tenant-unit show detailed` output to the default tenant-unit shown, if any, in the output.

```
# ddboost user option set ddbu1 default-tenant-unit tu1
```

6. Create a new alert notification group with the same name as the alert notification group in the `Alerts` section of the `smt tenant-unit show detailed` output.

```
# alert notify-list create tu1_grp tenant-unit tu1
```

7. Assign each email address in the alert notification group in the `Alerts` section of the `smt tenant-unit show detailed` output to the new alert notification group.

```
# alert notify-list add tu1_grp emails tom.tenant@abc.com
```

24

DD Secure Multitenancy

This chapter includes:

Topics:

Secure Multitenancy overview

Secure Multitenancy (SMT) is the simultaneous hosting, by an internal IT department or an external provider, of an IT infrastructure for more than one consumer or workload (business unit, department, or Tenant).

SMT provides the ability to securely isolate many users and workloads in a shared infrastructure, so that the activities of one Tenant are not apparent or visible to the other Tenants.

A *Tenant* is a consumer (business unit, department, or customer) who maintains a persistent presence in a hosted environment.

Within an enterprise, a Tenant may consist of one or more business units or departments on a protection system that is configured and managed by IT staff.

- For a business unit (BU) use case, the Finance and Human Resources departments of a corporation could share the same system, but each department would be unaware of the presence of the other.
- For a service provider (SP) use case, the SP could deploy one or more systems to accommodate different Protection Storage services for multiple end-customers.

Both use cases emphasize the separation of different customer data on the same physical system.

SMT architecture basics

Secure Multitenancy (SMT) provides a simple approach to setting up Tenants and Tenant Units, using MTrees. SMT setup is performed using DD Management Center and/or the DDOS command line interface. This administration guide provides the theory of SMT and some general command line instructions.

The basic architecture of SMT is as follows.

- A Tenant is created on the DD Management Center and/or DD system.
- A Tenant Unit is created on a DD system for the Tenant.
- One or more MTrees are created to meet the storage requirements for the Tenant's various types of backups.
- The newly created MTrees are added to the Tenant Unit.
- Backup applications are configured to send each backup to its configured Tenant Unit MTree.

NOTE: For more information about DD Management Center, see the *DD Management Center User Guide*. For more information about the DDOS command line interface, see the *DDOS Command Reference*.

Terminology used in Secure Multitenancy (SMT)

Understanding the terminology that is used in SMT will help you better understand this unique environment.

MTrees

MTrees are logical partitions of the file system and offer the highest degree of management granularity, meaning users can perform operations on a specific MTree without affecting the entire file system. MTrees are assigned to Tenant Units and contain that Tenant Unit's individualized settings for managing and monitoring SMT.

Multi-Tenancy

Multi-Tenancy refers to the hosting of an IT infrastructure by an internal IT department, or an external service provider, for more than one consumer/workload (business unit/department/Tenant) simultaneously. DD SMT enables *Data Protection-as-a-Service*.

RBAC (role-based access control)

RBAC offers multiple roles with different privilege levels, which combine to provide the administrative isolation on a multi-tenant protection system.

Storage Unit

A *Storage Unit* is an MTree configured for the DD Boost protocol. Data isolation is achieved by creating a Storage Unit and assigning it to a DD Boost user. The DD Boost protocol permits access only to Storage Units assigned to DD Boost users connected to the system.

Tenant

A *Tenant* is a consumer (business unit/department/customer) who maintains a persistent presence in a hosted environment.

Tenant Self-Service

Tenant Self-Service is a method of letting a Tenant log in to a protection system to perform some basic services (view MTrees or storage units that belong to the tenant unit, or change the tenant's own password). This reduces the bottleneck of always having to go through an administrator for these basic tasks. The Tenant can access only their assigned Tenant Units. Tenant Users and Tenant Admins will, of course, have different privileges.

Tenant Unit

A *Tenant Unit* is the partition of a system that serves as the unit of administrative isolation between Tenants. Tenant units that are assigned to a tenant can be on the same or different systems and are secured and logically isolated from each other, which ensures security and isolation of the control path when running multiple Tenants simultaneously on the shared infrastructure. Tenant Units can contain one or more *MTrees*, which hold all configuration elements that are needed in a multi-tenancy setup. Users, management-groups, notification-groups, and other configuration elements are part of a Tenant Unit.

Control path and network isolation

Control path isolation is achieved by providing the user roles of *tenant-admin* and *tenant-user* for a Tenant Unit. *Network isolation* for data and administrative access is achieved by associating a fixed set of *data access IP address(es)* and *management IP address(es)* with a Tenant Unit.

The *tenant-admin* and *tenant-user* roles are restricted in scope and capability to specific Tenant Units and to a restricted set of operations they can perform on those Tenant Units. To ensure a logically secure and isolated data path, a system administrator

must configure one or more Tenant Unit MTrees for each protocol in an SMT environment. Supported protocols include DD Boost, NFS, CIFS, and DD VTL. Access is strictly regulated by the native access control mechanisms of each protocol.

Tenant-self-service sessions (through ssh) can be restricted to a fixed set of *management IP address(es)* on a DD system. Administrative access sessions (through ssh/http/https) can also be restricted to a fixed set of management IP address(es) on DD systems. By default, however, there are no management IP address(es) associated with a Tenant Unit, so the only standard restriction is through the use of the *tenant-admin* and *tenant-user* roles. You must use `smt tenant-unit management-ip` to add and maintain management IP address(es) for Tenant Units.

Similarly, data access and data flow (into and out of Tenant Units) can be restricted to a fixed set of local or remote *data access IP address(es)*. The use of assigned data access IP address(es) enhances the security of the DD Boost and NFS protocols by adding SMT-related security checks. For example, the list of storage units returned over DD Boost RPC can be limited to those which belong to the Tenant Unit with the assigned local data access IP address. For NFS, access and visibility of exports can be filtered based on the local data access IP address(es) configured. For example, using `showmount -e` from the local data access IP address of a Tenant Unit will only display NFS exports belonging to that Tenant Unit.

The *sysadmin* must use `smt tenant-unit data-ip` to add and maintain data access IP address(es) for Tenant Units.

NOTE: If you attempt to mount an MTree in an SMT using a non-SMT IP address, the operation will fail.

If multiple Tenant Units are belong to the same tenant, they can share a default gateway. However, if multiple Tenant Units that belong to different tenants are oprevented from using the same default gateway.

Multiple Tenant Units belonging to the same tenant can share a default gateway. Tenant Units that belong to different tenants cannot use the same default gateway.

Understanding RBAC in SMT

In Secure Multitenancy (SMT), permission to perform a task depends on the role that is assigned to a user. DDMC uses role-based access control (RBAC) to control these permissions.

All DDMC users can:

- View all tenants
- Create, read, update, or delete tenant units belonging to any tenant if the user is an administrator on the protection system hosting the tenant unit
- Assign and unassign tenant units to and from a tenant if the user is an administrator on the system hosting the tenant unit
- View tenant units belonging to any tenant if the user has any assigned role on the system hosting the tenant unit

To perform more advanced tasks depends on the role of the user, as follows:

admin role

A user with an *admin* role can perform all administrative operations on a protection system. An *admin* can also perform all SMT administrative operations on the system, including setting up SMT, assigning SMT user roles, enabling tenant self-service mode, creating a tenant, and so on. In the context of SMT, the *admin* is typically referred to as the *landlord*. In DDOS, the role is known as the *sysadmin*.

To have permission to edit or delete a tenant, you must be both a DDMC *admin* and a DDOS *sysadmin* on all systems that are associated with the tenant units of that tenant. If the tenant does not have any tenant units, you need only to be a DDMC *admin* to edit or delete that tenant.

limited-admin role

A user with a *limited-admin* role can perform all administrative operations on a system as the *admin*. However, users with the *limited-admin* role cannot delete or destroy MTrees. In DDOS, there is an equivalent *limited-admin* role.

tenant-admin role

A user with a *tenant-admin* role can perform certain tasks only when *tenant self-service* mode is enabled for a specific tenant unit. Responsibilities include scheduling and running a backup application for the tenant and monitoring resources and statistics within the assigned tenant unit. The *tenant-admin* can view audit logs, but RBAC ensures that only audit logs from the tenant

units belonging to the *tenant-admin* are accessible. In addition, *tenant-admins* ensure administrative separation when tenant self-service mode is enabled. In the context of SMT, the *tenant-admin* is referred to as the *backup admin*.

tenant-user role

A user with a *tenant-user* role can monitor the performance and usage of SMT components only on tenant unit(s) assigned to them and only when tenant self-service is enabled, but a user with this role cannot view audit logs for their assigned tenant units. Also, *tenant-users* may run the `show` and `list` commands.

none role

A user with a role of *none* is not allowed to perform any operations on a system other than changing their password and accessing data using DD Boost. However, after SMT is enabled, the *admin* can select a user with a *none* role from the system and assign them an SMT-specific role of *tenant-admin* or *tenant-user*. Then, that user can perform operations on SMT management objects.

management groups

BSPs (backup service providers) can use *management groups* defined in a single, external AD (active directory) or NIS (network information service) to simplify managing user roles on tenant units. Each BSP tenant may be a separate, external company and may use a name-service such as AD or NIS.

With SMT management groups, the AD and NIS servers are set up and configured by the *admin* in the same way as SMT local users. The *admin* can ask their AD or NIS administrator to create and populate the group. The *admin* then assigns an SMT role to the entire group. Any user within the group who logs in to the system is logged in with the role that is assigned to the group.

When users leave or join a tenant company, they can be removed or added to the group by the AD or NIS administrator. It is not necessary to modify the RBAC configuration on a system when users who are part of the group are added or removed.

Provisioning a Tenant Unit

Launching the configuration wizard begins the initial provisioning procedure for Secure Multitenancy (SMT). During the procedure, the wizard creates and provisions a new Tenant Unit based on Tenant configuration requirements. Information is entered by the administrator, as prompted. After completing the procedure, the administrator proceeds to the next set of tasks, beginning with enabling Tenant Self-Service mode. Following the initial setup, manual procedures and configuration modifications can be performed as required.

Steps

1. Start SMT.

```
# smt enable
SMT enabled.
```

2. Verify that SMT is enabled.

```
# smt status
SMT is enabled.
```

3. Launch the SMT configuration wizard.

```
# smt tenant-unit setup
No tenant-units.
```

4. Follow the configuration prompts.

```
SMT TENANT-UNIT Configuration

  Configure SMT TENANT-UNIT at this time (yes|no) [no]: yes

  Do you want to create new tenant-unit (yes/no)? : yes

Tenant-unit Name
  Enter tenant-unit name to be created
```

```
  : SMT_5.7_tenant_unit
Invalid tenant-unit name.
  Enter tenant-unit name to be created
  : SMT_57_tenant_unit

Pending Tenant-unit Settings
Create Tenant-unit   SMT_57_tenant_unit

Do you want to save these settings (Save|Cancel|Retry): save
SMT Tenant-unit Name Configurations saved.

SMT TENANT-UNIT MANAGEMENT-IP Configuration

  Configure SMT TENANT-UNIT MANAGEMENT-IP at this time (yes|no) [no]: yes

  Do you want to add a local management ip to this tenant-unit? (yes|no) [no]: yes

port  enabled  state  DHCP        IP address              netmask     type additional
                                                       /prefix length        setting
----- ------- ------- ---- -------------------------- -------------- ---- ----------
ethMa   yes   running  no  192.168.10.57              255.255.255.0   n/a
                           fe80::260:16ff:fe49:f4b0** /64
eth3a   yes   running ipv4 192.168.10.236*            255.255.255.0*  n/a
                           fe80::260:48ff:fe1c:60fc** /64
eth3b   yes   running  no  192.168.50.57              255.255.255.0   n/a
                           fe80::260:48ff:fe1c:60fd** /64
eth4b   yes   running  no  192.168.60.57              255.255.255.0   n/a
                           fe80::260:48ff:fe1f:5183** /64
----- ------- ------- ---- -------------------------- -------------- ---- ----------
* Value from DHCP
** auto_generated IPv6 address

Choose an ip from above table or enter a new ip address. New ip addresses will need
to be created manually.

Ip Address
  Enter the local management ip address to be added to this tenant-unit
  : 192.168.10.57

  Do you want to add a remote management ip to this tenant-unit? (yes|no) [no]:

Pending Management-ip Settings

Add Local Management-ip    192.168.10.57
  Do you want to save these settings (Save|Cancel|Retry): yes
  unrecognized input, expecting one of Save|Cancel|Retry

  Do you want to save these settings (Save|Cancel|Retry): save
  Local management access ip "192.168.10.57" added to tenant-unit "SMT_57_tenant_unit".

SMT Tenant-unit Management-IP Configurations saved.

SMT TENANT-UNIT MANAGEMENT-IP Configuration

  Do you want to add another local management ip to this tenant-unit? (yes|no) [no]:

  Do you want to add another remote management ip to this tenant-unit? (yes|no) [no]:

SMT TENANT-UNIT DDBOOST Configuration
  Configure SMT TENANT-UNIT DDBOOST at this time (yes|no) [no]:

SMT TENANT-UNIT MTREE Configuration
  Configure SMT TENANT-UNIT MTREE at this time (yes|no) [no]: yes

Name                       Pre-Comp (GiB)   Status   Tenant-Unit
-------------------------  --------------   ------   -----------
/data/col1/laptop_backup          4846.2   RO/RD     -
/data/col1/random                23469.9   RO/RD     -
/data/col1/software2              2003.7   RO/RD     -
/data/col1/tsm6                 763704.9   RO/RD     -
-------------------------  --------------   ------   -----------
D : Deleted
Q : Quota Defined
```

```
RO : Read Only
RW : Read Write
RD : Replication Destination
IRH : Retention-Lock Indefinite Retention Hold Enabled
ARL : Automatic-Retention-Lock Enabled
RLGE : Retention-Lock Governance Enabled
RLGD : Retention-Lock Governance Disabled
RLCE : Retention-Lock Compliance Enabled


  Do you want to assign an existing MTree to this tenant-unit? (yes|no) [no]:

  Do you want to create a mtree for this tenant-unit now? (yes|no) [no]: yes

MTree Name
  Enter MTree name
  :  SMT_57_tenant_unit
Invalid mtree path name.
  Enter MTree name
  : SMT_57_tenant_unit


Invalid mtree path name.
  Enter MTree name
  : /data/col1/SMT_57_tenant_unit

MTree Soft-Quota
  Enter the quota soft-limit to be set on this MTree (<n> {MiB|GiB|TiB|PiB}|none)
  :

MTree Hard-Quota
  Enter the quota hard-limit to be set on this MTree (<n> {MiB|GiB|TiB|PiB}|none)
  :

Pending MTree Settings
Create MTree        /data/col1/SMT_57_tenant_unit
MTree Soft Limit    none
MTree Hard Limit    none

  Do you want to save these settings (Save|Cancel|Retry): save
  MTree "/data/col1/SMT_57_tenant_unit" created successfully.
  MTree "/data/col1/SMT_57_tenant_unit" assigned to tenant-unit  "SMT_57_tenant_unit".

SMT Tenant-unit MTree Configurations saved.

SMT TENANT-UNIT MTREE Configuration

Name                        Pre-Comp (GiB)   Status   Tenant-Unit
-------------------------   --------------   ------   -----------
/data/col1/laptop_backup           4846.2   RO/RD     -
/data/col1/random                 23469.9   RO/RD     -
/data/col1/software2               2003.7   RO/RD     -
/data/col1/tsm6                  763704.9   RO/RD     -
-------------------------   --------------   ------   -----------
D : Deleted
Q : Quota Defined
RO : Read Only
RW : Read Write
RD : Replication Destination
IRH : Retention-Lock Indefinite Retention Hold Enabled
ARL : Automatic-Retention-Lock Enabled
RLGE : Retention-Lock Governance Enabled
RLGD : Retention-Lock Governance Disabled
RLCE : Retention-Lock Compliance Enabled


  Do you want to assign another MTree to this tenant-unit? (yes|no) [no]: yes

  Do you want to assign an existing MTree to this tenant-unit? (yes|no) [no]:

  Do you want to create another mtree for this tenant-unit? (yes|no) [no]:

SMT TENANT-UNIT SELF-SERVICE Configuration
```

```
  Configure SMT TENANT-UNIT SELF-SERVICE at this time (yes|no) [no]: yes
  Self-service of this tenant-unit is disabled

  Do you want to enable self-service of this tenant-unit? (yes|no) [no]: yes

  Do you want to configure a management user for this tenant-unit? (yes|no) [no]:

  Do you want to configure a management group for this tenant-unit (yes|no) [no]: yes

Management-Group Name
  Enter the group name to be assigned to this tenant-unit
  : SMT_57_tenant_unit_group

What role do you want to assign to this group  (tenant-user|tenant-admin) [tenant-user]:
tenant-admin

Management-Group Type
  What type do you want to assign to this group (nis|active-directory)?
  : nis

Pending Self-Service Settings
Enable Self-Service       SMT_57_tenant_unit
Assign Management-group   SMT_57_tenant_unit_group
Management-group role     tenant-admin
Management-group type     nis

  Do you want to save these settings (Save|Cancel|Retry): save
  Tenant self-service enabled for tenant-unit "SMT_57_tenant_unit"
  Management group "SMT_57_tenant_unit_group" with type "nis" is assigned to tenant-unit
  "SMT_57_tenant_unit" as "tenant-admin".

SMT Tenant-unit Self-Service Configurations saved.

SMT TENANT-UNIT SELF-SERVICE Configuration

  Do you want to configure another management user for this tenant-unit? (yes|no) [no]:

  Do you want to configure another management group for this tenant-unit? (yes|no) [no]:

SMT TENANT-UNIT ALERT Configuration

  Configure SMT TENANT-UNIT ALERT at this time (yes|no) [no]: yes
  No notification lists.

Alert Configuration

Alert Group Name
  Specify alert notify-list group name to be created
  : SMT_57_tenant_unit_notify

Alert email addresses
  Enter email address to receive alert for this tenant-unit
  : dd_proserv@emc.com

  Do you want to add more emails (yes/no)?
  : no

Pending Alert Settings
Create Notify-list group   SMT_57_tenant_unit_notify
Add emails                 dd_proserv@emc.com

  Do you want to save these settings (Save|Cancel|Retry): save
  Created notification list "SMT_57_tenant_unit_notify" for tenant "SMT_57_tenant_unit".
  Added emails to notification list "SMT_57_tenant_unit_notify".

SMT Tenant-unit Alert Configurations saved.

Configuration complete.
```

Enabling Tenant Self-Service mode

For administrative separation of duties and delegation of administrative/management tasks to implement Tenant Self-Service, which is required for control path isolation, the system administrator can enable this mode on a Tenant Unit and then assign users to manage the unit in the roles of tenant-admin or tenant-user. These roles allow users other than the administrator to perform specific tasks on the Tenant Unit to which they are assigned. In addition to administrative separation, Tenant Self-Service mode helps reduce the management burden on internal IT and service provider staff.

Steps

1. View Tenant Self-Service mode status for one or all Tenant Units.

```
# smt tenant-unit option show { tenant-unit | all }
```

2. Enable Tenant Self-Service mode on the selected Tenant Unit.

```
# smt tenant-unit option set tenant-unit self-service { enabled | disabled }
```

Data access by protocol

Secure data paths, with protocol-specific access controls, enable security and isolation for Tenant Units. In a Secure Multitenancy (SMT) environment, data access protocol management commands are also enhanced with a Tenant Unit parameter to enable consolidated reporting.

DD systems support multiple data access protocols simultaneously, including DD Boost, NFS, CIFS, and DD VTL. A DD system can present itself as an application-specific interface, such as a file server offering NFS or CIFS access over the Ethernet, a DD VTL device, or a DD Boost device.

The native access control mechanisms of each supported protocol ensure that the data paths for each Tenant remain separate and isolated. Such mechanisms include access control lists (ACLs) for CIFS, exports for NFS, DD Boost credentials, and Multi-User Boost credential-aware access control.

Multi-User DD Boost and Storage Units in SMT

Storage Unit ownership sets the user permissions for Multi-User DD Boost with Secure Multitenancy (SMT).

Multi-User DD Boost describes using multiple DD Boost user credentials for DD Boost Access Control, in which each user has a separate username and password.

A *Storage Unit* is an MTree configured for the DD Boost protocol. A user can be associated with, or "own," one or more Storage Units. A user cannot own Storage Units that another user owns. Only the user owning the Storage Unit can access the Storage Unit for any type of data access, such as backup or restore. The number of DD Boost usernames cannot exceed the maximum number of MTrees. Storage Units that are associated with SMT must have the *none* role that is assigned to them.

Each backup application must authenticate by using its DD Boost username and password. After authentication, DD Boost verifies the authenticated credentials to confirm ownership of the Storage Unit. The backup application is granted access to the Storage Unit. This action occurs only if the backup application user credentials match the usernames that are associated with the Storage Unit. If user credentials and usernames do not match, the job fails with a permission error.

Configuring access for CIFS

Common Internet File System (CIFS) is a file-sharing protocol for remote file access. In a Secure Multitenancy (SMT) configuration, backup and restores require client access to the CIFS shares residing in the MTree of the associated Tenant Unit. Data isolation is achieved using CIFS shares and CIFS ACLs.

Steps

1. Create an MTree for CIFS and assign the MTree to the tenant unit.

```
# mtree create mtree-path tenant-unit tenant-unit
```

2. Set capacity soft and hard quotas for the MTree.

```
# mtree create mtree-path tenant-unit tenant-unit ] [quota-soft-limit n {MiB|GiB|TiB|
PiB} ] [quota-hard-limit n {MiB|GiB|TiB|PiB}
```

3. Create a CIFS share for *pathname* from the MTree.

```
# cifs share create share path pathname clients clients
```

Configuring NFS access

NFS is a UNIX-based, file-sharing protocol for remote file access. In a Secure Multitenancy (SMT) environment, backup and restores require client access to the NFS exports residing in the MTree of the associated Tenant Unit. Data isolation is achieved using NFS exports and network isolation. NFS determines if an MTree is associated with a network-isolated Tenant Unit. If so, NFS verifies the connection properties associated with the Tenant Unit. Connection properties include the destination IP address and interface or client hostname.

Steps

1. Create an MTree for NFS and assign the MTree to the tenant unit.

```
# mtree create mtree-path tenant-unit tenant-unit
```

2. Set capacity soft and hard quotas for the MTree.

```
# mtree create mtree-path tenant-unit tenant-unit ] [quota-soft-limit n {MiB|GiB|TiB|
PiB} ] [quota-hard-limit n {MiB|GiB|TiB|PiB}
```

3. Create an NFS export by adding one or more clients to the MTree.

```
# nfs add path client-list
```

Configuring access for DD VTL

DD VTL Tenant data isolation is achieved using DD VTL access groups that create a virtual access path between a host system and the DD VTL. (The physical Fibre Channel connection between the host system and DD VTL must already exist.)

Placing tapes in the DD VTL allows them to be written to, and read by, the backup application on the host system. DD VTL tapes are created in a DD VTL pool, which is an MTree. Because DD VTL pools are MTrees, the pools can be assigned to Tenant Units. This association enables SMT monitoring and reporting.

For example, if a tenant-admin is assigned a Tenant Unit that contains a DD VTL pool, the tenant-admin can run MTree commands to display read-only information. Commands can run only on the DD VTL pool assigned to the Tenant Unit.

These commands include:

- `mtree list` to view a list of MTrees in the Tenant Unit
- `mtree show compression` to view statistics on MTree compression
- `mtree show performance` to view statistics on performance

Output from most `list` and `show` commands include statistics that enable service providers to measure space usage and calculate chargeback fees.

DD VTL operations are unaffected and continue to function normally.

Using DD VTL NDMP TapeServer

DD VTL Tenant data isolation is also achieved using NDMP. DDOS implements a NDMP (Network Data Management Protocol) tape server that allows NDMP-capable systems to send backup data to the DD system via a three-way NDMP backup.

The backup data is written to virtual tapes (which are in a pool) by a DD VTL assigned to the special DD VTL group *TapeServer*.

Because the backup data is written to tapes in a pool, information in the DD VTL topic regarding MTrees also applies to the DD NDMP TapeServer.

Data management operations

Secure Multitenancy (SMT) management operations include monitoring Tenant Units and other objects, such as Storage Units and MTrees. For some SMT objects, additional configuration or modification may also be required.

Collecting performance statistics

Each MTree can be measured for performance or "usage" statistics and other real-time information. Historical consumption rates are available for DD Boost Storage Units. Command output lets the tenant-admin collect usage statistics and compression ratios for an MTree associated with a Tenant Unit, or for all MTrees and associated Tenant Units. Output may be filtered to display usage in intervals ranging from minutes to months. Results are passed to the administrator, who uses the statistics as a chargeback metric. A similar method is used to gather usage statistics and compression ratios for Storage Units.

Steps

1. Collect MTree real-time performance statistics.

   ```
   # mtree show performance
   ```

2. Collect performance statistics for MTrees associated with a Tenant Unit.

   ```
   # mtree show performance tenant-unit <tenant-unit-name>
   ```

3. Collect compression statistics for MTrees associated with a Tenant Unit.

   ```
   # mtree show compression
   ```

Modifying quotas

To meet QoS criteria, a system administrator uses DDOS "knobs" to adjust the settings required by the Tenant configuration. For example, the administrator can set "soft" and "hard" quota limits on DD Boost Storage Units. Stream "soft" and "hard" quota limits can be allocated only to DD Boost Storage Units assigned to Tenant Units. After the administrator sets the quotas, the tenant-admin can monitor one or all Tenant Units to ensure no single object exceeds its allocated quotas and deprives others of system resources.

About this task

Quotas are set initially when prompted by the configuration wizard, but they can be adjusted or modified later. The example below shows how to modify quotas for DD Boost. (You can also use `quota capacity` and `quota streams` to deal with capacity and stream quotas and limits.)

Steps

1. To modify soft and hard quota limits on DD Boost Storage Unit "su33":
   `ddboost storage-unit modify su33 quota-soft-limit 10 Gib quota-hard-limit 20 Gib`
2. To modify stream soft and hard limits on DD Boost Storage Unit "su33":
   `ddboost storage-unit modify su33 write-stream-soft-limit 20 read-stream-soft-limit 6 repl -stream-soft-limit 20 combined-stream-soft-limit 20`
3. To report physical size for DD Boost Storage Unit "su33":
   `ddboost storage-unit modify su33 report-physical-size 8 GiB`

SMT and replication

In case of disaster, user roles dictate how a user can assist in data recovery operations. Several replication types are available in an SMT configuration. (See the *DD Replicator* chapter for more detail on how to perform replication.)

Here are some points to consider regarding user roles:

- The admin can recover MTrees from a replicated copy.
- The tenant-admin can replicate MTrees from one system to another, using DD Boost managed file replication.
- The tenant-admin can recover MTrees from a replicated copy, also by using DD Boost managed file replication.

Collection replication

Collection replication replicates core Tenant Unit configuration information.

Secure replication over public internet

To protect against man-in-the-middle (MITM) attacks when replicating over a public internet connection, authentication includes validating SSL certificate-related information at the replication source and destination.

MTree replication (NFS/CIFS)

MTree replication is supported on MTrees assigned to Tenant Units. During MTree replication, an MTree assigned to a Tenant Unit on one system can be replicated to an MTree assigned to a Tenant Unit on another system. MTree replication is not allowed between two different Tenants on the two DD systems. When security mode is set to *strict*, MTree replication is allowed only when the MTrees belong to same Tenants.

For backward compatibility, MTree replication from an MTree assigned to a Tenant Unit to an unassigned MTree is supported, but must be configured manually. Manual configuration ensures the destination MTree has the correct settings for the Tenant Unit. Conversely, MTree replication from an unassigned MTree to an MTree assigned to a Tenant Unit is also supported.

When setting up SMT-aware MTree replication, *security mode* defines how much checking is done on the Tenant. The *default* mode checks that the source and destination do not belong to different Tenants. The *strict* mode makes sure the source and destination belong to the same Tenant. Therefore, when you use strict mode, you must create a Tenant on the destination machine with the same UUID as the UUID of the Tenant on the source machine that is associated with the MTree being replicated.

DD Boost managed file replication (also with DD Boost AIR)

DD Boost managed file replication is supported between Storage Units, regardless of whether one Storage Unit, or both, are assigned to Tenant Units.

During DD Boost managed file replication, Storage Units are not replicated in total. Instead, certain files within a Storage Unit are selected by the backup application for replication. The files selected in a Storage Unit and assigned to a Tenant Unit on one system can be replicated to a Storage Unit assigned to a Tenant Unit on another system.

For backward compatibility, selected files in a Storage Unit assigned to a Tenant Unit can be replicated to an unassigned Storage Unit. Conversely, selected files in an unassigned Storage Unit can be replicated to a Storage Unit assigned to a Tenant Unit.

DD Boost managed file replication can also be used in DD Boost AIR deployments.

Replication control for QoS

An upper limit on replication throughput (`repl-in`) can be specified for an MTree. Since MTrees for each tenant are assigned to a Tenant Unit, each tenant's replication resource usage can be capped by applying these limits. The relation of this feature to SMT is that MTree Replication is subject to this throughput limit.

SMT Tenant alerts

A DD system generates *events* when it encounters potential problems with software or hardware. When an event is generated, an *alert* notification is sent immediately via email to members designated in the notification list and to the system administrator.

SMT alerts are specific to each Tenant Unit and differ from DD system alerts. When Tenant Self-Service mode is enabled, the tenant-admin can choose to receive alerts about the various system objects he or she is associated with and any critical events, such as an unexpected system shutdown. A tenant-admin may only view or modify notification lists to which he or she is associated.

The example below shows a sample alert. Notice that the two event messages at the bottom of the notification are specific to a Multi-Tenant environment (indicated by the word "Tenant"). For the entire list of DDOS and SMT alerts, see the *DDOS MIB Quick Reference Guide* or the SNMP MIB.

```
EVT-ENVIRONMENT-00021 - Description: The system has been shutdown by abnormal method; for
example, not by one of the following: 1) Via IPMI chassis control command 2) Via power
button 3) Via OS shutdown.

Action: This alert is expected after loss of AC (main power) event. If this shutdown is
not expected and persists, contact your contracted support provider or visit us online at
https://dell.com/support.

Tenant description: The system has experienced an unexpected power loss and has restarted.

Tenant action: This alert is generated when the system restarts after a power loss. If
this alert repeats, contact your System Administrator.
```

Managing snapshots

A *snapshot* is a read-only copy of an MTree captured at a specific point in time. A snapshot can be used for many things, for example, as a restore point in case of a system malfunction. The required role for using `snapshot` is admin or tenant-admin.

To view snapshot information for an MTree or a Tenant Unit:

```
# snapshot list mtree mtree-path | tenant-unit tenant-unit
```

To view a snapshot schedule for an MTree or a Tenant Unit:

```
# snapshot schedule show [name | mtrees mtree-listmtree-list | tenant-unit tenant-unit]
```

Performing a file system Fast Copy

A Fast Copy operation clones files and directory trees of a source directory to a target directory on a DD system. There are special circumstances regarding Fast Copy with Secure Multitenancy (SMT).

Here are some considerations when performing a file system Fast Copy with Tenant Self-Service mode enabled:

- A tenant-admin can Fast Copy files from one Tenant Unit to another when the tenant-admin is the tenant-admin for both Tenant Units, and the two Tenant Units belong to the same Tenant.
- A tenant-admin can Fast Copy files within the same Tenant Unit.
- A tenant-admin can Fast Copy files within the Tenant Units at source and destination.

To perform a file system Fast Copy:

```
# filesys fastcopy source <src> destination <dest>
```

25

Cloud Tier

This chapter includes:

Topics:

Cloud Tier overview

Cloud Tier is a native feature of DDOS 6.0 (or later) for moving data from the active tier to low-cost, high-capacity object storage in the public, private, or hybrid cloud for long-term retention. Cloud Tier is best suited for long-term storage of infrequently accessed data that is being held for compliance, regulatory, and governance reasons. The ideal data for Cloud Tier is data that is past its normal recovery window.

Cloud Tier is managed using a single protection system namespace. There is no separate cloud gateway or virtual appliance required. Data movement is supported by the native policy management framework. Conceptually, the cloud storage is treated as an additional storage tier (Cloud Tier) attached to the system, and data is moved between tiers as needed. File system metadata associated with the data stored in the cloud is maintained in local storage, and also mirrored to the cloud. The metadata that resides in local storage facilitates operations such as deduplication, cleaning, Fast Copy, and replication. This local storage is divided into self-contained buckets, called cloud units, for ease of manageability.

Supported platforms

Cloud Tier is supported on physical platforms that have the necessary memory, CPU, and storage connectivity to accommodate another storage tier.

Cloud Tier is supported on these systems:

Table 97. Cloud Tier supported configurations

Model	Memory	Cloud capacity	Required number of SAS I/O modules	Supported disk shelf types for metadata storage	Number of ES30 shelves, ES40 shelves, DS60 disk packs, disk packs required	Required capacity for metadata storage
DD3300 4 TB	16 GB	8 TB	N/A	N/A	N/A	1 x 1 TB virtual disk = 1 TB
DD3300 8 TB	48 GB	16 TB	N/A	N/A	N/A	2 x 1 TB virtual disks = 2 TB

Table 97. Cloud Tier supported configurations (continued)

Model	Memory	Cloud capacity	Required number of SAS I/O modules	Supported disk shelf types for metadata storage	Number of ES30 shelves, ES40 shelves, DS60 disk packs, disk packs required	Required capacity for metadata storage
DDD3300 16 TB	48 GB	32 TB	N/A	N/A	N/A	2 x 1 TB virtual disks = 2 TB
DD3300 32 TB	64 GB	64 TB	N/A	N/A	N/A	4 x 1 TB virtual disks = 4 TB
DD6400 (default)	192 GB	64 TB	N/A	N/A	N/A	Cloud Tier metadata is stored on the 8 TB drives in the DD6400 controller.
DD6400 (expanded)	192 GB	344 TB	1	ES40	N/A	Cloud Tier metadata is stored on the 8 TB drives in the DD6400 controller and ES40 shelves.
DD6800	192 GB	576 TB	2	DS60 or ES30	2	30 x 4 TB HDDs = 120 TB
DD6900	288 GB	576 TB	2	DS60, ES40, or ES30[a]	2	30 x 4 TB HDDs = 120 TB
DD9300	384 GB	1400 TB	2	DS60 or ES30	4	60 x 4 TB HDDs = 240 TB
DD9400	576 GB	1536 TB	2	DS60, ES40, or ES30[a]	4	60 x 4 TB HDDs = 240 TB
DD9800	768 GB	2016 TB	4	DS60 or ES30	5	75 x 4 TB HDDs = 300 TB
DD9900 (default)	1152 GB	2016 TB	2	DS60, ES40, or ES30[a]	5	75 x 4 TB HDDs = 300 TB
DD9900 (expanded)	1152 GB	3072 TB	2	DS60, ES40, or ES30[a]	6	90 x 4 TB HDDs = 360 TB
DDVE 16 TB	32 GB	32 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB[b]
DDVE 64 TB	60 GB	128 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB[b]
DDVE 96 TB	80 GB	192 TB	N/A	N/A	N/A	1 x 500 GB virtual disk = 500 GB[b]

a. ES30 shelves are only supported after a controller upgrade from an older system model.
b. The minimum metadata size is a hard limit. Dell recommends that you start with 1 TB for metadata storage and expand in 1 TB increments. The *DDVE Installation and Administration Guide* provides more details about using Cloud Tier with DDVE.

NOTE: Cloud Tier is not supported on any system that is not listed and is configured with Collection Replication.

NOTE: The Cloud Tier feature may consume all available bandwidth in a shared WAN link, especially in a low bandwidth configuration (1 Gbps), and this may impact other applications sharing the WAN link. If there are shared applications on the WAN, the use of QoS or other network limiting is recommended to avoid congestion and ensure consistent performance over time.

If bandwidth is constrained, the rate of data movement will be slow and you will not be able to move as much data to the cloud. It is best to use a dedicated link for data going to the Cloud Tier.

NOTE: Do not send traffic over onboard management network interface controllers (ethMx interfaces).

NOTE: Separate Cloud Tier licenses are not required to configure both cloud units on the system. A single license is distributed between both cloud units. Actual Cloud Tier storage capacity may be lower than the aggregated total of both cloud units post-comp, and the licensed Cloud Tier capacity.

Cloud Tier performance

The system uses internal optimizations to maximize Cloud Tier performance.

Cloud seeding

The current migration engine to cloud is filed based and an efficient de-duplication optimized engine is used for identifying and migrating only unique segments to cloud. This file based migration engine's efficiency is high when migrating higher generation data to Cloud Tier, which already has some data to de-duplicate against. However, when Cloud Tier is empty or nearly empty, there is no data to de-duplicate against. There is an overhead of compute cycles that are invested in deduplication. With seeding-based migration, the dededuplication filtering is maintained on active tier storage and only unique data is migrated in bulk to Cloud Tier. In cloud seeding, the engine migrates the content from local storage to cloud storage without processing it for deduplication. When cloud seeding is active, files that are marked for migration to cloud storage are not cleaned (i.e. space is not freed-up) as part of the active tier file system cleaning until the migration of all identified files by seeding is complete. Active tier storage must be sized to account for this in environments where large amounts of data are migrated to cloud storage. If the Cloud Tier storage is less than five percent full and has post-comp data usage of 30 TiB (or more), as seen in `filesys show space` command, the system automatically uses cloud seeding when migrating data to cloud storage.

After five percent of the Cloud Tier capacity is consumed, cloud seeding automatically deactivates. Data is then processed for deduplication before migration to cloud storage.

Here are additional points to consider when using Seeding migration:

- Migration is supported in Seeding mode only when:
  - Active tier postcomp used size is 30 TiB or more as reported in filesys show space output.
  - Active tier is less than 70% full, when migration starts as reported in filesys show space output.

    NOTE: While in seeding mode, if Active Tier usage during a migration cycle exceeds 90%, migration is halted and restarted in regular Filecopy mode.
- Migration in seeding mode is auto-suspended by cleaning on active tier, for the entire duration of the active tier cleaning. Once cleaning completes, seeding resumes automatically and restarts the migration to cloud.
- Migration in seeding mode auto-suspends if a cloud UNAVAIL event is received on the cloud-unit (cloud-unit is reported as "disconnected") to which it is migrating, and only resumes once the cloud-unit is available reports as active.
- Cleaning cannot start on a cloud-unit that is the destination of an in-progress migration operation in Seeding mode.

  NOTE: In two cloud-unit systems, to force cleaning to start on a second cloud-unit which is not being seeded, suspend migration in seeding mode using the `data-movement suspend` command and run the `cloud clean start` command on the second cloud-unit.
- Probabilistic File Verification in cloud does not run against cloud-units on which seeding mode migration is in progress, even if it is the default policy.
- If cleaning is in progress on Active Tier or Cloud Tier and scheduled data movement starts in seeding mode, the data movement operation suspends for the duration of the cleaning activity.
- If data movement is stopped or suspended while in seeding mode, the `data-movement status` command will not show an updated total for the number of files moved once the data movement operation is complete, resulting in a mismatch between the `Files eligible` and `Files moved` values. If the `Files failed` value is zero, the data movement operation completed successfully.

- Migration in seeding mode does not migrate files from MTrees which are replication destinations, even if the files are eligible for migration. Files from these replication destination MTrees are migrated with the Filecopy engine once migration in seeding mode from all eligible MTree is complete.
- Seeding mode migration suspends physical capacity reporting for the duration of the migration activity.
- Migration in Seeding mode is only supported on all cloud enabled systems and configurations that have more than 80 Gb of RAM. Seeding based migration is disabled by default for DD VEs.

Large object size

Cloud Tier uses object sizes of 1 MB or 4 MB (depending on the cloud storage provider) to reduce the metadata overhead, and lower the number of objects to migrate to cloud storage.

Resuming scheduled cleaning on Cloud Tier storage

DDOS versions 7.4 and higher provide the ability to automatically suspend and resume scheduled Cloud Tier cleaning operations under certain conditions. To maintain system performance, the suspend-resume feature is only supported on systems with a minimum of 96 GB of memory. No user intervention is required.

Cloud cleaning is a scheduled activity that runs as function of the cleaning frequency set for the Active Tier. The frequency determines the interval between two cleaning operations. However, depending on the amount of data to clean, the cleaning operation can take a lot of time, which makes it prone to environmental and scheduled interruptions (such as scheduled cleaning on Active Tier storage or network connectivity to Cloud Tier storage), causing it to abort the current run.

If the cleaning operation has reached the copy phase, Scheduled cleaning activity on Cloud Tier storage is now protected against environmental and scheduled interruptions by entering a suspended state when the interruption begins and resuming once the interruption is resolved or completed. Scheduled cleaning on Cloud Tier storage aborts on the invocation of the `filesys clean start` command, or if the file system restarts. Cleaning also aborts if an HA failover occurs while the cleaning operation is suspended.

NOTE: The suspend-resume functionality does not work on cleaning operations initiated with the `cloud clean start` command.

Data movement to cloud storage while cloud cleaning is suspended is supported in environments with one cloud unit, if the cloud unit is active. The cleaning operation can resume while the data movement operation is in progress. However, in environments with two cloud units, if data movement starts on cloud unit 1 while cleaning started and was suspended on cloud unit 2, the cleaning operation remains suspended until the data movement operation is complete.

NOTE: If the data movement operation starts in seeding mode (Cloud Seeding), the suspended cleaning operation aborts.

The suspend-resume feature tracks file system capacity usage between two internal watermarks:

- Low watermark: Suspend-resume for scheduled cloud cleaning operations is only supported when Active Tier capacity is less than 80%. Beyond 80%, scheduled cloud cleaning operations will not suspend.
- High water mark: If Active Tier capacity usage exceeds 90% while the cloud cleaning operation is suspended, the cloud cleaning operating detects this condition on resume and aborts the current cleaning run.

Configuring Cloud Tier

To configure Cloud Tier, add the license and enclosures, set a system passphrase, and create a file system with support for data movement to the cloud.

- For Cloud Tier, the cloud capacity license is required to configure cloud units and move data to the cloud.
- To license Cloud Tier, refer to the applicable *DDOS Release Notes* for the most up-to-date information on product features, software updates, software compatibility guides, and information about protection products, licensing, and service.
- To set a system passphrase, use the **Administration** > **Access** > **Administrator Access** tab.

  If the system passphrase is not set, the **Set Passphrase** button appears in the Passphrase area. If a system passphrase is configured, the **Change Passphrase** button appears, and your only option is to change the passphrase.

- To create a file system, use the File System Create Wizard.

Configuring storage for Cloud Tier

Cloud Tier storage is required for the DD system to support cloud-units. The Cloud Tier holds the metadata for the migrated files, while the actual data resides in the cloud.

Prerequisites

The file system must be disabled to configure Cloud Tier.

Steps

1. Select **Data Management** > **File System** and click **Disable** (at the bottom of the screen) to disable the file system.
2. Select **Hardware** > **Storage**.
3. In the Overview tab, expand **Cloud Tier**.
4. Click **Configure**.
   The Configure Cloud Tier dialog box is displayed.
5. Select the checkbox for the shelf to be added from the Addable Storage section.

   CAUTION: DD3300 systems require the use of 1 TB storage devices for Cloud Tier metadata storage.

6. Click the **Add to Tier** button.
7. Click **Save** to add the storage.
8. Select **Data Management** > **File System** and click **Enable Cloud Tier**.
   To enable the cloud tier, you must meet the storage requirement for the licensed capacity. Configure the cloud tier of the file system. Click **Next**.

   A cloud file system requires a local store for a local copy of the cloud metadata.

9. Click **Enable**.
   The cloud tier is enabled with the designated storage.
10. Click **OK**.
    You must create cloud units separately, after the file system is enabled.
11. Select **Enable file system**.

Cleanable Space Estimation

The Cleanable Space Estimation tool assesses the amount of space that can be reclaimed on the Active Tier when the data-movement process migrates eligible files to the cloud and GC cleans the file system.

This tool can work with or without a cloud license present.

When there is no active CLOUDTIER-CAPACITY license, manually provide the age-threshold to use to assess total cleanable space on the active tier. If there is both an age-threshold and there is a policy set on MTrees, the preference is given to the user provided age-threshold.

There are three workflows:

- A system with cloud migration policies set: Files are identified as "eligible" based on the policy set on the respective MTrees and calculates the cleanable space.
- A system with cloud migration policies set but with a user provided age-threshold: Files are identified based on the user given age-threshold, overriding the system policies.
- A system with no cloud: Mandatory requirement for user to provide an age-threshold which would be used to determine total cleanable space.

Some additional points to consider:

- Data-movement cannot run in parallel with the data-movement eligibility-check process.
- Cleaning on Active Tier cannot be started if the eligibility-check is running.
- The eligibility-check cannot start if cleaning on the Active Tier is running.
- Cleaning on Cloud Tier cannot be started if the eligibility-check is running.
- The eligibility-check cannot start if cleaning on the Cloud Tier is running.
- If an UNAVAIL event is received, it should not have any impact on the eligibility-check operation.

- If the file system stops or crashes, eligibility-check stops and does not auto-resume once file system comes back up again.

NOTE: There is no provision for initiating the eligibility-check from DD System Manager.

Configuring cloud units

The cloud tier consists of a maximum of two cloud units, and each cloud unit is mapped to a cloud provider, enabling multiple cloud providers per protection system. The system must be connected to the cloud and have an account with a supported cloud provider.

Configuring cloud units includes these steps:

- Configuring the network, including firewall and proxy settings
- Importing CA certificates
- Adding cloud units

Firewall and proxy settings

Network firewall ports

- Port 443 (HTTPS) and/or Port 80 (HTTP) must be open to the cloud provider networks for both the endpoint IP and the provider authentication IP for bi-directional traffic.

  For example, for Amazon S3, both s3.ap-southeast-1.amazonaws.com and s3.amazonaws.com must have port 80 and/or port 443 unblocked and set to allow bi-directional IP traffic.

  NOTE: Several public cloud providers use IP ranges for their endpoint and authentication addresses. In this situation, the IP ranges used by the provider need to be unblocked to accommodate potential IP changes.
- Remote cloud provider destination IP and access authentication IP address ranges must be allowed through the firewall.
- For ECS private cloud, local ECS authentication and web storage (S3) access IP ranges and ports 9020 (HTTP) and 9021 (HTTPS) must be allowed through local firewalls.

  NOTE: ECS private cloud load balancer IP access and port rules must also be configured.

Proxy settings

If there are any existing proxy settings that cause data above a certain size to be rejected, those settings must be changed to allow object sizes up to 4.5MB.

If customer traffic is being routed through a proxy, the self-signed/CA-signed proxy certificate must be imported. See "Importing CA certificates" for details.

OpenSSL cipher suites

- Ciphers - ECDHE-RSA-AES256-SHA384, AES256-GCM-SHA384
- TLS Version: 1.2

NOTE: Default communication with all cloud providers is initiated with strong cipher.

Supported protocols

- HTTP
- HTTPS

NOTE: Default communication with all public cloud providers occurs on secure HTTP (HTTPS), but you can overwrite the default setting to use HTTP.

Importing CA certificates

Before you can add cloud units for Alibaba, Amazon Web Services S3 (AWS), Azure, Elastic Cloud Storage (ECS), and Google Cloud Provider (GCP), you must import CA certificates.

Prerequisites

For AWS and Azure public cloud providers, root CA certificates can be downloaded from https://www.digicert.com/digicert-root-certificates.htm.

- For an AWS cloud provider, download the Baltimore CyberTrust Root certificate and the Starfield Class 2 Certification Authority.
- For an Azure cloud provider, download the Baltimore CyberTrust Root certificate.
- For ECS, the root certificate authority varies by customer.

  Implementing cloud storage on ECS requires a load balancer. If an HTTPS endpoint is used as an endpoint in the configuration, be sure to import the root CA certificate. Contact your load balancer provider for details.

- For an S3 Flexible provider, import the root CA certificate. Contact your S3 Flexible provider for details.

If your downloaded certificate has a .crt extension, it is likely that it will need to be converted to a PEM-encoded certificate. If so, use OpenSSL to convert the file from .crt format to .pem (for example, `openssl x509 -inform der -in BaltimoreCyberTrustRoot.crt -out BaltimoreCyberTrustRoot.pem`).

- For Alibaba:
  1. Download the GlobalSign Root R1 certificate from https://support.globalsign.com/customer/portal/articles/1426602-globalsign-rootcertificates.
  2. Convert the downloaded certificate to a PEM-encoded format. The OpenSSL command for this conversion is: `openssl x509 -inform der -in <root_cert.crt> -out <root_cert.pem>`.
  3. Import the certificate to the system.
- For GCP:
  1. Download the GTS Root R1 certificate from https://pki.goog:
     a. Click **Repository**.
     b. Scroll down to **Download CA certificates**.
     c. Expand **Root CAs**.
     d. Click **Action**.
     e. Download the certificate.
  2. Convert the downloaded certificate to a PEM-encoded format. The OpenSSL command for this conversion is: `openssl x509 -inform der -in <root_cert.crt> -out <root_cert.pem>`.
  3. Import the certificate to the system.

About this task

If the cloud provider changes the certificate, the cloud units enter a disconnected state. The system detects this change and triggers an alert for the administrator to import the new certificates to restore access.

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. In the tool bar, click **Manage Certificates**.
   The Manage Certificates for Cloud dialog is displayed.
3. Click **Add**.
4. Select one of these options:
   - **I want to upload the certificate as a .pem file.**

     Browse to and select the certificate file.

   - **I want to copy and paste the certificate text.**
     - Copy the contents of the .pem file to your copy buffer.
     - Paste the buffer into the dialog.
5. Click **Add**.

Next steps

If the cloud provider changes the certificate, the cloud units enter a disconnected state. The system detects this change and triggers an alert for the administrator to import the new certificates to restore access.

Adding a cloud unit for Elastic Cloud Storage (ECS)

About this task

A protection system or DDVE instance requires a close time synchronization with the ECS system to configure a DD cloud unit. Configuring NTP on the protection system or DDVE instance, and the ECS system addresses this issue.

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.
   The **Add Cloud Unit** dialog box appears.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
   The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Elastic Cloud Storage (ECS)** from the list.
5. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
6. Enter the provider **Access key** as password text.

   NOTE: Use the ECS username as the access key.

7. Enter the provider **Secret key** as password text.
8. Enter the provider **Endpoint** in this format: **`http://<ip/hostname>:<port>`**. If you are using a secure endpoint, use **`https`** instead.

   NOTE: Implementing cloud storage on ECS requires a load balancer.

   By default, ECS runs the S3 protocol on port 9020 for HTTP and 9021 for HTTPS. With a load balancer, these ports are sometimes remapped to 80 for HTTP and 443 for HTTPS, respectively. Check with your network administrator for the correct ports.
9. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
   Enter the proxy hostname, port, user, and password.

   NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.

10. Click **Add**.
    The File System main window displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Alibaba

About this task

Regions are configured at bucket level instead of object level. Therefore, all objects contained in a bucket are stored in the same region. A region is specified when a bucket is created, and cannot be changed once it is created.

Table 98. Alibaba regions

Regions	Location	Region Name
Mainland China regions	China East 1 (Hangzhou)	oss-cn-hangzhou
	China East 2 (Shanghai)	oss-cn-shanghai
	China North 1 (Qingdao)	oss-cn-qingdao

Table 98. Alibaba regions (continued)

Regions	Location	Region Name
	China North 2 (Beijing)	oss-cn-beijing
	China North 3 (zhangjiakou)	oss-cn-zhangjiakou
	China North 5 (huhehaote)	oss-cn-huhehaote
	China South 1 (Shenzhen)	oss-cn-shenzhen
International Regions	Hong Kong	oss-cn-hongkong
	US West 1 (Silicon Valley)	oss-us-west-1
	US East 1 (Virginia)	oss-us-east-1
	Asia Pacific SE 1 (Singapore)	oss-ap-southeast-1
	Asia Pacific SE 2 (Sydney)	oss-ap-southeast-2
	Asia Pacific SE 3 (Kuala Lumpur)	oss-ap-southeast-3
	Asia Pacific SE 5 (Jakarta)	oss-ap-southeast-5
	Asia Pacific NE 1 (Tokyo)	oss-ap-northeast-1
	Asia Pacific SOU 1 (Mumbai)	oss-ap-south-1
	EU Central 1 (Frankfurt)	oss-eu-central-1
	Middle East 1 (Dubai)	oss-me-east-1

The Alibaba Cloud user credentials must have permissions to create and delete buckets and to add, modify, and delete files within the buckets they create. AliyunOSSFullAccess is preferred, but these are the minimum requirements:

- ListBuckets
- GetBucket
- PutBucket
- DeleteBucket
- GetObject
- PutObject
- DeleteObject

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.
   The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
   The remaining fields in the **Add Cloud Unit** dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Alibaba Cloud** from the drop-down list.
5. Select **Standard** or **IA** from the **Storage class** drop-down list.
6. Select the region from the **Storage region** drop-down list.
7. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
8. Enter the provider **Access key** as password text.
9. Enter the provider **Secret key** as password text.
10. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the Alibaba cloud provider occurs on port 443.
11. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
    Enter the proxy hostname, port, user, and password.

    NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.

12. Click **Add**.

The file system main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Amazon Web Services S3

AWS offers a range of storage classes. The *Cloud Providers Compatibility Matrix*, available from E-Lab Navigator provides up-to-date information about the supported storage classes.

About this task

For enhanced security, the Cloud Tier feature uses Signature Version 4 for all AWS requests. Signature Version 4 signing is enabled by default.

The following endpoints are used by the AWS cloud provider, depending on storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

FIPS-compliant endpoints are available for AWS Government Cloud.

Starting in DDOS 7.8, the us-east-1 region no longer supports the legacy endpoint s3.amazonaws.com. The us-east-1 region now requires the endpoint s3.us-east-1.amazonaws.com. Verify the firewall is open to reach the new endpoint before upgrading to DDOS 7.8.

- s3.us-east-1.amazonaws.com
- s3.us-east-2.amazonaws.com
- s3.us-west-1.amazonaws.com
- s3.us-west-2.amazonaws.com
- s3.eu-west-1.amazonaws.com
- s3.af-south-1.amazonaws.com
- s3.ap-east-1.amazonaws.com
- s3.ap-northeast-1.amazonaws.com
- s3.ap-northeast-2.amazonaws.com
- s3.ap-northeast-3.amazonaws.com
- s3.ap-southeast-1.amazonaws.com
- s3.ap-southeast-2.amazonaws.com
- s3.ap-southeast-3.amazonaws.com
- s3.sa-east-1.amazonaws.com
- s3.ap-south-1.amazonaws.com
- s3.eu-central-1.amazonaws.com
- s3.eu-north-1.amazonaws.com
- s3.eu-south-1.amazonaws.com
- s3.eu-west-2.amazonaws.com
- s3.eu-west-3.amazonaws.com
- s3.us-gov-east-1.amazonaws.com
- s3.fips.us-gov-east-1.amazonaws.com
- s3.us-gov-west-1.amazonaws.com
- s3.fips.us-gov-west-1.amazonaws.com
- s3.ca-central-1.amazonaws.com
- s3.me-central-1.amazonaws.com
- s3.me-south-1.amazonaws.com

NOTE: The China region is not supported.

NOTE: The AWS user credentials must have permissions to create and delete buckets and to add, modify, and delete files within the buckets they create. S3FullAccess is preferred, but these are the minimum requirements:

- CreateBucket
- ListBucket
- DeleteBucket
- ListAllMyBuckets
- GetObject
- PutObject

- DeleteObject

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.
   The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
   The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Amazon Web Services S3** from the drop-down list.
5. Select the storage class from the drop-down list.
6. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
7. Select the appropriate **Storage region** from the drop-down list.
8. Enter the provider **Access key** as password text.
9. Enter the provider **Secret key** as password text.
10. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the AWS cloud provider occurs on port 443.
11. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
    Enter the proxy hostname, port, user, and password.

    NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.

12. Click **Add**.
    The file system main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Azure

Microsoft Azure offers a range of storage account types. The *Cloud Providers Compatibility Matrix*, available from E-Lab Navigator provides up-to-date information about the supported storage classes.

About this task

The following endpoints are used by the Azure cloud provider, depending on account type. Be sure that DNS can resolve these hostnames before configuring cloud units.

- For Azure Public accounts: *<account-name>*.blob.core.windows.net

  NOTE: Do not include the domain **`blob.core.windows.net`** as part of the account name.
- For Azure Government accounts: *<account-name>*.blob.core.usgovcloudapi.net

  NOTE: Do not include the domain **`blob.core.usgovcloudapi.net`** as part of the account name.
- For Azure China accounts: *<account-name>*.blob.core.chinacloudapi.cn

  NOTE: Do not include the domain **`blob.core.chinacloudapi.cn`** as part of the account name.

The account name is obtained from the Azure cloud provider console.

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.
   The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.
   The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Microsoft Azure Storage** from the drop-down list.
5. For **Account type**, select **Government**, **Public**, or **China**.
6. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
7. Enter the provider **Account name**.

8. Enter the provider **Primary key** as password text.
9. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the Azure cloud provider occurs on port 443.
10. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
    Enter the proxy hostname, port, user, and password.

    NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.

11. Click **Add**.
    The file system main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding a cloud unit for Google Cloud Provider

About this task

The following tables list the Cloud Storage locations available for storing data.

Table 99. Multi-regional locations

Multi-regional name	Multi-regional description
Asia	Data centers in Asia
US	Data centers in the United States
EU	Data centers in the European Union

Table 100. Regional locations

Regional locations	Location	Region name
North America	northamerica-northeast1	Montréal
	us-central1	Iowa
	us-east1	South Carolina
	us-east4	Northern Virginia
	us-west1	Oregon
South America	southamerica-east1	São Paulo
Europe	europe-north1	Finland
	europe-west1	Belgium
	europe-west2	London
	europe-west3	Frankfurt
	europe-west4	Netherlands
Asia	asia-east1	Taiwan
	asia-northeast1	Tokyo
	asia-south1	Mumbai
	asia-southeast1	Singapore
Australia	australia-southeast1	Sydney

The Google Cloud Provider user credentials must have permissions to create and delete buckets and to add, modify, and delete files within the buckets they create. These are the minimum requirements:

- ListBucket
- PutBucket
- GetBucket

- DeleteBucket
- GetObject
- PutObject
- DeleteObject

NOTE:

Cloud Tier only supports Nearline and is selected automatically during setup.

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.

   The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.

   The remaining fields in the **Add Cloud Unit** dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Google Cloud Storage** from the drop-down list.
5. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
6. Enter the provider **Access key** as password text.
7. Enter the provider **Secret key** as password text.
8. **Storage class** is set as **Nearline** by default.

   If a multi-regional location is selected (Asia, EU or US), then the storage class and the location constraint is Nearline Multi-regional. All other regional locations have the storage class set as Nearline Regional.
9. Select the **Region**.
10. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with Google Cloud Provider occurs on port 443.
11. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.

    Enter the proxy hostname, port, user, and password.

    NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.
12. Click **Add**.

    The file system main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Adding an S3 Flexible provider cloud unit

The Cloud Tier feature supports additional qualified S3 cloud providers under an S3 Flexible provider configuration option. The *Cloud Providers Compatibility Matrix*, available from E-Lab Navigator provides up-to-date information about the supported S3 cloud providers.

About this task

The S3 Flexible provider option supports the standard and standard-infrequent-access storage classes. The endpoints will vary depending on cloud provider, storage class and region. Be sure that DNS is able to resolve these hostnames before configuring cloud units.

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click **Add**.

   The Add Cloud Unit dialog is displayed.
3. Enter a name for this cloud unit. Only alphanumeric characters are allowed.

   The remaining fields in the Add Cloud Unit dialog pertain to the cloud provider account.
4. For **Cloud provider**, select **Flexible Cloud Tier Provider Framework for S3** from the drop-down list.
5. In the **Bucket** field, optionally specify a pre-existing, empty bucket to use the for the cloud unit.
6. Enter the provider **Access key** as password text.
7. Enter the provider **Secret key** as password text.

8. Specify the appropriate **Storage region**.
9. Enter the provider **Endpoint** in this format: **`http://<ip/hostname>:<port>`**. If you are using a secure endpoint, use **`https`** instead.
10. For **Storage class**, select the appropriate storage class from the drop-down list.
11. Ensure that port 443 (HTTPS) is not blocked in firewalls. Communication with the S3 cloud provider occurs on port 443.
12. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
    Enter the proxy hostname, port, user, and password.

    NOTE: There is an optional step to run the cloud provider verify tool before adding the cloud unit. This tool performs pre-check tests to ensure that all requirements are met before to adding the actual cloud unit.

13. Click **Add**.
    The File System main window now displays summary information for the new cloud unit as well a control for enabling and disabling the cloud unit.

Modifying a cloud unit or cloud profile

About this task

Modify cloud unit credentials, an S3 Flexible provider name, or details of a cloud profile.

Modifying cloud unit credentials

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click the pencil icon for the cloud unit whose credentials you want to modify.
   The Modify Cloud Unit dialog is displayed.
3. For **Account name**, enter the new account name.
4. For **Access key**, enter the new provider access key as password text.

   NOTE: Modifying the access key is not supported for ECS environments.

5. For **Secret key**, enter the new provider secret key as password text.
6. For **Primary key**, enter the new provider primary key as password text.

   NOTE: Modifying the primary key is only supported for Azure environments.

7. If an HTTP proxy server is required to get around a firewall for this provider, click **Configure** for **HTTP Proxy Server**.
8. Click **OK**.

Modifying an S3 Flexible provider name

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.
2. Click the pencil icon for the S3 Flexible cloud unit whose name you want to modify.
   The Modify Cloud Unit dialog is displayed.
3. For **S3 Provider Name**, enter the new provider name.
4. Click **OK**.

Modifying the AWS proxy configuration

Steps

1. Select **Data Management** > **File System** > **Cloud Units**.

2. Click the pencil icon for the AWS cloud unit you want to modify.
   The Modify Cloud Unit dialog is displayed.
3. For **HTTP Proxy Server**, click **Configure** to specify new configuration values, or click **Clear** to remove the proxy server configuration.
4. Click **Save**.

Using the CLI to modify a cloud profile

Steps

Run the `cloud profile modify` command to modify the details of a cloud profile. The system prompts you to modify individual details of the cloud profile.

For AWS S3, run this command to add a storage class to an existing cloud profile.

The profile details that can be modified depend on the cloud provider:

- Alibaba Cloud supports modification of the access key, and secret key.
- The AWS S3 storage class can only be modified from:

  NOTE: Data that already exists in cloud storage are not changed when the storage class changes, use AWS to change the storage class of this data. All new data added to cloud storage after the change are in the new storage class.

  - Standard to Standard-Infrequent Access
  - Standard to Glacier IR
  - Standard-Infrequent Access to Glacier IR
- Azure supports modification of the primary key.
- ECS supports modification of the secret key.
- S3 Flexible supports modification of the access key, secret key, and provider name.

Deleting a cloud unit

This operation results in the loss of all data in the cloud unit selected for deletion. Be sure to delete all files before deleting the cloud units.

Prerequisites

- Check if a security policy is configured on the system. Cloud units cannot be deleted without security officer authorization.
- Check if data movement to the cloud is running (CLI command: data-movement status). If it is, stop data movement using the “data-movement stop” CLI command.
- Check if cloud cleaning is running for this cloud unit (CLI command: cloud clean status). If it is, stop cloud cleaning using the “cloud clean” CLI command.
- Check if a data movement policy is configured for this cloud unit (CLI command: data-movement policy show). If it is, remove this policy using the “data-movement policy reset” CLI command.

Steps

1. Use the following CLI command to identify files in the cloud unit.

   ```
   # filesys report generate file-location
   ```

2. Delete the files that are in the cloud unit to be deleted.
3. Use the following CLI command to run cloud cleaning.

   ```
   # cloud clean start unit-name
   ```

   Wait for cleaning to complete. The cleaning may take time depending on how much data is present in the cloud unit.
4. Disable the file system.
5. Use the following CLI command to delete the cloud unit.

   NOTE: The `cloud unit del` command requires security officer authorization, and cannot be run if a security policy is not configured on the system. If multifactor authentication is enabled, the security officer must enter an RSA SecurID token after their password.

```
# cloud unit del unit-name
```

Internally, this marks the cloud unit as DELETE_PENDING.

6. Use the following CLI command to validate that the cloud unit is in the DELETE_PENDING state.

```
# cloud unit list
```

7. Enable the file system.
   The file system initiates the procedure in the background to delete any remaining objects from the buckets in the cloud for this cloud unit and then delete the buckets. This process can take a long time, depending on how many objects were remaining in these buckets. Until the bucket cleanup completes, this cloud unit continues to consume a slot on the protection system, which may prevent creation of a new cloud unit if both slots are occupied.
8. Periodically check the state using this CLI command:

```
# cloud unit list
```

   The state remains DELETE_PENDING while the background cleanup is running.
9. Verify from the cloud provider S3 portal that all corresponding buckets have been deleted and the associated space has been freed up.
10. If needed, reconfigure data movement policies for affected MTrees and restart data movement.

Results

If you have difficulty completing this procedure, contact Support.

Data movement

Data is moved from the active tier to the cloud tier as specified by your individual data movement policy. The policy is set on a per-MTree basis. Data movement can be initiated manually or automatically using a schedule.

Adding data movement policies to MTrees

A file is moved from the Active to the Cloud Tier based on the date it was last modified. For data integrity, the entire file is moved at this time. The *Data Movement Policy* establishes the file age threshold, age range, and the destination.

About this task

NOTE: A data movement policy cannot be configured for the /backup MTree.

Steps

1. Select **Data Management** > **MTree**.
2. In the top panel, select the MTree to which you want to add a data movement policy.
3. Click the **Summary** tab.
4. Under **Data Movement Policy** click **Add**.
5. For **File Age in Days**, set the file age threshold (**Older than**) and optionally, the age range (**Younger than**).

   NOTE: The minimum number of days for **Older than** is 14. For nonintegrated backup applications, files moved to the cloud tier cannot be accessed directly and need to be recalled to the active tier before you can access them. So, choose the age threshold value as appropriate to minimize or avoid the need to access a file moved to the cloud tier.
6. For **Destination**, specify the destination cloud unit.
7. Click **Add**.

Next steps

NOTE: If a data movement policy is not configured, scheduled Cloud Tier backups for PowerProtect Data Manager will fail.

Moving data manually

You can start and stop data movement manually. Any MTree that has a valid data movement policy has its files moved.

Steps

1. Select **Data Management** > **File System**.
2. At the bottom of the page, click **Show Status of File System Services**.
   These status items are displayed:
   - File System
   - Physical Capacity Measurement
   - Data Movement
   - Active Tier Cleaning
3. For **Data Movement**, click **Start**.

Moving data automatically

You can move data automatically, using a schedule and a throttle. Schedules can be daily, weekly, or monthly.

Steps

1. Select **Data Management** > **File System** > **Settings**.
2. Click the **Data Movement** tab.
3. Set the throttle and schedule.

   NOTE: The throttle is for adjusting resources for internal system processes; it does not affect network bandwidth.

   NOTE: If a cloud unit is inaccessible when cloud tier data movement runs, the cloud unit is skipped in that run. Data movement on that cloud unit occurs in the next run if the cloud unit becomes available. The data movement schedule determines the duration between two runs. If the cloud unit becomes available and you cannot wait for the next scheduled run, you can start data movement manually.

Recalling a file from the Cloud Tier

For nonintegrated backup applications, you must recall the data to the active tier before you can restore the data. Backup administrators must trigger a recall or backup applications must perform a recall before cloud-based backups can be restored. Once a file is recalled, its aging is reset and starts again from 0, and the file is eligible based on the age policy set. A file can be recalled on the same MTree only. Integrated applications can restore a file directly.

About this task

NOTE: In an MTree replication context, the file is read-only on the destination MTree.

NOTE: If a file resides only in a snapshot, it cannot be recalled directly. To recall a file in a snapshot, use fastcopy to copy the file from the snapshot back to the active MTree, then recall the file from the cloud. A file can only be recalled from the cloud to an active MTree.

Recall operations are subject to the following limitations:

- The recall operation will fail if insufficient space is available in Active Tier storage.
- MTree level recalls can run for one MTree at a time.
- If an MTree recall is in progress, no new file or MTree recalls can be started.
- Data movement to cloud storage and MTree recall operations cannot be run at the same time.

File-level and MTree-level recall operations restart automatically if there is a cloud unavailable event or file system panic.

Steps

1. Select **Data Management** > **File System** > **Summary**.

2. Do one of the following:
   - In the Cloud Tier section of the Space Usage panel, click **Recall**.
   - Expand the File System status panel at the bottom of the screen and click **Recall**.

   NOTE: The **Recall** link is available only if a cloud unit is created and has data.
3. In the Recall File from Cloud dialog, specify individual files or an MTree to recall: Click **Recall**.
   - To recall individual files, select **Recall files by name** and enter the exact file name (no wildcards) and full path of the file to be recalled, for example: `/data/col1/mt11/file1.txt`.
   - To recall an Mtree, select select **Recall all files for an MTree** and enter the MTree path, for example: `/data/col1/mt11`.
4. Click **Recall**.
5. To check the status of the recall, do one of the following:
   - In the Cloud Tier section of the Space Usage panel, click **Details**.
   - Expand the File System status panel at the bottom of the screen and click **Details**.

   The Cloud File Recall Details dialog is displayed, showing the file path or MTree name, cloud provider, recall progress, and amount of data transferred. If there are unrecoverable errors during the recall, an error message is displayed. Hover the cursor over the error message to display a tool tip with more details and possible corrective actions.

Results

Once the file has been recalled to the active tier, you can restore the data.

NOTE: For nonintegrated applications, once a file has been recalled from the cloud tier to the active tier, a minimum of 14 days must elapse before the file is eligible for data movement. After 14 days, normal data movement processing will occur for the file. The file now has to wait the age-threshold or age-range to move back to the cloud as this time the ptime will be examined rather than the mtime. This restriction does not apply to integrated applications.

NOTE: For data-movement, nonintegrated applications configure an age-based data movement policy on the protection system to specify which files get migrated to the cloud tier, and this policy applies uniformly to all files in an MTree. Integrated applications use an application-managed data movement policy, which lets you identify specific files to be migrated to the cloud tier.

Using the CLI to recall a file from the cloud tier

For nonintegrated backup applications, you must recall the data to the active tier before you can restore the data. Backup administrators must trigger a recall or backup applications must perform a recall before cloud-based backups can be restored. Once a file is recalled, its aging is reset and will start again from 0, and the file will be eligible based on the age policy set. A file can be recalled on the source MTree only. Integrated applications can recall a file directly.

About this task

NOTE: If a file resides only in a snapshot, it cannot be recalled directly. To recall a file in a snapshot, use fastcopy to copy the file from the snapshot back to the active MTree, then recall the file from the cloud. A file can only be recalled from the cloud to an active MTree.

Perform a file-level recall

Steps

1. Check the location of the file using: `filesys report generate file-location [path {<path-name> | all}] [output-file <filename>] [tiering]`

   The pathname can be a file or directory; if it is a directory, all files in the directory are listed.

```
Filename              Location
--------              --------
/data/col1/mt11/file1.txt   Cloud Unit 1
```

2. Recall the file using: `data-movement recall path <path-name>`

   This command is asynchronous, and it starts the recall.

```
data-movement recall path /data/col1/mt11/file1.txt
Recall started for "/data/col1/mt11/file1.txt".
```

3. Monitor the status of the recall using `data-movement status [path {pathname | all | [queued] [running] [completed] [failed]} | to-tier cloud | all]`

```
# data-movement status path /data/col1/stu-av1/FILE-115113.0005.0006

Data-movement recall:
---------------------
Path Name (*)                                  Status       File Size       Logical
Logical   Time Queued (**)    Time Started        Time Ended
                                                                        Bytes Moved   Bytes
Verified
---------------------------------------   ---------   ----------   -----------
-------------   -----------------   -----------------   -----------------
/data/col1/stu-av1/FILE-115113.0005.0006   running       1.00 GiB      1.00 GiB
22.46 MiB   Jul 31 2020 06:29   Jul 31 2020 06:29   -
---------------------------------------   ---------   ----------   -----------
-------------   -----------------   -----------------   -----------------
(*)  The completed or failed recall jobs prior to the last filesystem restart were not
displayed.
(**) The queued timestamp may not be available if the filesystem was restarted.
```

If the status shows that the recall isn't running for a given path, the recall may have finished, or it may have failed.

4. Verify the location of the file using `filesys report generate file-location [path {<path-name> | all}] [output-file <filename>] [tiering]`

```
Filename             Location
--------                      --------
/data/col1/mt11/file1.txt   Active
```

Results

Once the file has been recalled to the Active Tier, you can restore the data.

NOTE: For nonintegrated applications, once a file has been recalled from the Cloud Tier to the Active Tier, a minimum of 14 days must elapse before the file is eligible for data movement. After 14 days, normal data movement processing will occur for the file. This restriction does not apply to integrated applications.

NOTE: For data-movement, nonintegrated applications configure an age-based data movement policy on the protection system to specify which files get migrated to the cloud tier, and this policy applies uniformly to all files in an MTree. Integrated applications use an application-managed data movement policy, which lets you identify specific files to be migrated to the cloud tier.

Perform an MTree-level recall

Steps

1. Check the location of the data using: `filesys report generate file-location [path {<path-name> | all}] [output-file <filename>] [tiering]`

   Identify the MTree.

```
Filename             Location
--------                      --------
/data/col1/mt11/file1.txt   Cloud Unit 1
```

2. Recall the file using: `data-movement recall MTree <mtree-path>`

   This command is asynchronous, and it starts the recall.

```
data-movement recall mtree /data/col1/mt11

        ** Mtree recall can potentially consume significant space from Active Tier and
may incur some data retrieval cost.
        You can choose to recall specific files using 'data-movement recall path'
```

```
command.
        Do you want to continue? (yes|no) [no]: yes

Recall initiated for "/data/col1/mt11". Run the "data-movement recall watch" command to
monitor its progress.
```

3. Monitor the status of the recall using `data-movement recall status detailed`

```
# data-movement recall status detailed

Data-movement to active tier:
-----------------------------
Data-movement recall:
80% complete; Elapsed time: 0:00:31
Moved (post-comp): 94.08 MiB, (pre-comp): 1000.00 MiB,
Files inspected: 10, Files eligible: 10, Files moved: 1, Files failed: 0

Data-movement recall:
---------------------
Path Name  Status File Size Logical Logical Time Queued (**) Time Started Time Ended
Bytes Moved Bytes Verified
------------------------------ ------- ---------- ----------- --------------
----------------- ----------------- ----------
/data/col1/stu1/data.0006.0000 running 200.00 MiB 200.00 MiB - Dec 14 2020 03:26 Dec 14
2020 03:26 -
/data/col1/stu1/data.0003.0000 running 200.00 MiB 200.00 MiB - Dec 14 2020 03:26 Dec 14
2020 03:26 -
/data/col1/stu1/data.0002.0000 running 200.00 MiB 200.00 MiB - Dec 14 2020 03:26 Dec 14
2020 03:26 -
/data/col1/stu1/data.0001.0000 running 200.00 MiB 200.00 MiB - Dec 14 2020 03:26 Dec 14
2020 03:26 -
------------------------------ ------- ---------- ----------- --------------
----------------- ----------------- ----------
 The completed or failed recall jobs prior to the last filesystem restart were not
displayed.
(**) The queued timestamp may not be available if the filesystem was restarted.
Use 'data-movement status path all' option to list all recall jobs (up to the maximum
supported limit).
```

If the status shows that the recall isn't running for a given MTree, the recall may have finished, or it may have failed.

4. Verify the location of the data using `filesys report generate file-location [path {<path-name> | all}] [output-file <filename>] [tiering]`

```
Filename             Location
--------                      --------
/data/col1/mt11/file1.txt   Active
```

Results

Once the MTree has been recalled to the Active Tier, you can restore the data.

NOTE: For nonintegrated applications, once a file has been recalled from the Cloud Tier to the Active Tier, a minimum of 14 days must elapse before the file is eligible for data movement. After 14 days, normal data movement processing will occur for the file. This restriction does not apply to integrated applications.

NOTE: For data-movement, nonintegrated applications configure an age-based data movement policy on the protection system to specify which files get migrated to the cloud tier, and this policy applies uniformly to all files in an MTree. Integrated applications use an application-managed data movement policy, which lets you identify specific files to be migrated to the cloud tier.

Direct restore from the cloud tier

Direct restore lets nonintegrated applications read files directly from the Cloud Tier without going through the Active Tier.

Key considerations in choosing to use direct restore include:

- Direct restore does not require an integrated application and is transparent for nonintegrated applications.

- Reading from the cloud tier does not require copying first into the active tier.
- Histograms and statistics are available for tracking direct reads from the cloud tier.
- Direct restore is supported only for AWS and ECS cloud providers.
- Applications do experience cloud tier latency.
- Reading directly from the cloud tier is not bandwidth optimized.

The maximum number of files that can be recalled at one time depends on the system memory configuration:

- Systems with 256 GB of memory or more can recall up to 16 files at one time.
- Systems with less than 256 GB of memory can recall up to 8 files at one time.
- DD VE instances can recall up to 4 files at one time.

Direct restore is useful with nonintegrated applications that do not need to know about the cloud tier and won't need to restore cloud files frequently.

Using the CLI to configure Cloud Tier

You can use the CLI to configure Cloud Tier.

Steps

1. Configure storage for both active and cloud tier. As a prerequisite, the appropriate capacity licenses for both the active and cloud tiers must be installed.
   a. Ensure licenses for the features CLOUDTIER-CAPACITY and CAPACITY-ACTIVE are installed. To check the ELMS license:

      ```
      # elicense show
      ```

      If the license is not installed, use the `elicense update` command to install the license. Enter the command and paste the contents of the license file after this prompt. After pasting, ensure there is a carriage return, then press **Control-D** to save. You are prompted to replace licenses, and after answering yes, the licenses are applied and displayed.

      ```
      # elicense update
      Enter the content of license file and then press Control-D, or press Control-C to
      cancel.
      ```

   b. Display available storage:

      ```
      # storage show all
      # disk show state
      ```

   c. Add storage to the active tier:

      ```
      # storage add enclosures <enclosure no> tier active
      ```

   d. Add storage to the cloud tier:

      ```
      # storage add enclosures <enclosure no> tier cloud
      ```

2. Install certificates.

   Before you can create a cloud profile, you must install the associated certificates.

   For AWS and Azure public cloud providers, root CA certificates can be downloaded from https://www.digicert.com/digicert-root-certificates.htm.

   - For an AWS cloud provider, download the Baltimore CyberTrust Root certificate and the Starfield Class 2 Certification Authority.
   - For an Azure cloud provider, download the Baltimore CyberTrust Root certificate.
   - For Alibaba, download the GlobalSign Root R1 certificate from https://support.globalsign.com/customer/portal/articles/1426602-globalsign-rootcertificates.
   - For ECS, the root certificate authority will vary by customer. Contact your load balancer provider for details.

   Downloaded certificate files have a .crt extension. Use openssl on any Linux or Unix system where it is installed to convert the file from .crt format to .pem.

```
$openssl x509 -inform der -in DigiCertHighAssuranceEVRootCA.crt -out
DigiCertHighAssuranceEVRootCA.pem

$openssl x509 -inform der -in BaltimoreCyberTrustRoot.crt -out
BaltimoreCyberTrustRoot.pem

# adminaccess certificate import ca application cloud
Enter the certificate and then press Control-D, or press Control-C to cancel.
```

3. To configure the system for data-movement to the cloud, you must first enable the “cloud” feature and set the system passphrase if it has not already been set.

```
# cloud enable
Cloud feature requires that passphrase be set on the system.
Enter new passphrase:
Re-enter new passphrase:
Passphrases matched.
The passphrase is set.
Encryption is recommended on the cloud tier.
    Do you want to enable encryption? (yes|no) [yes]:
Encryption feature is enabled on the cloud tier.
Cloud feature is enabled.
```

4. Configure the cloud profile using the cloud provider credentials. The prompts and variables vary by provider.

```
# cloud profile add <profilename>
```

NOTE: For security reasons, this command does not display the access/secret keys you enter.

Select the provider:

```
Enter provider name (alibabacloud|aws|azure|ecs|google|s3_flexible)
```

- Alibaba Cloud requires access key, secret key, storage class and region.
- AWS S3 requires access key, secret key, storage class, and region.
- Azure requires account name, account type (Public, Government, or China), and primary key.
- ECS requires entry of access key, secret key and endpoint.
- Google Cloud Platform requires access key, secret key, and region. (Storage class is Nearline.)
- S3 Flexible providers require the provider name, access key, secret key, region, endpoint, and storage class.

At the end of each profile addition you are asked if you want to set up a proxy. If you do, these values are required: *proxy hostname*, *proxy port*, *proxy username*, and *proxy password*.

5. Verify the cloud profile configuration:

```
# cloud profile show
```

6. Create the active tier file system if it is not already created:

```
# filesys create
```

7. Enable the file system:

```
# filesys enable
```

8. Configure the cloud unit:

```
# cloud unit add unit-name profile profile-name [[bucket bucket-name] | [for-migration]]
```

Use the `cloud unit list` command to list the cloud units.

9. Optionally, configure encryption for the cloud unit.
   a. Verify that the ENCRYPTION license is installed:

```
# elicense show
```

   b. Enable encryption for the cloud unit:

```
# filesys encryption enable cloud-unit unitname
```

c. Check encryption status:

```
# filesys encryption status
```

10. Create one or more MTrees:

```
# mtree create /data/col1/mt11
```

11. Verify the Cloud Tier configuration:

```
# cloud provider verify
This operation will perform test data movement after creating a temporary profile and
bucket.
Do you want to continue? (yes|no) [yes]:
Enter provider name (alibabacloud|aws|azure|ecs|google|s3_flexible): aws
Enter the access key:
Enter the secret key:
Enter the storage class (STANDARD|STANDARD_IA|ONEZONE_IA|GLACIER_IR) [STANDARD]:
Enter the region (us-east-2|us-east-1|us-west-1|us-west-2|af-south-1|
                          ap-east-1|ap-southeast-3|ap-south-1|ap-northeast-3|
                          ap-northeast-2|ap-southeast-1|ap-southeast-2|
                          ap-northeast-1|ca-central-1|eu-central-1|eu-west-1|
                          eu-west-2|eu-south-1|eu-west-3|eu-north-1|me-south-1|
                     me-central-1|sa-east-1|us-gov-east-1|us-gov-west-1): us-west-1

Verifying cloud provider …
This process may take a few minutes.
Cloud Enablement Check:
 Checking Cloud feature enabled: PASSED
 Checking Cloud volume: PASSED

Connectivity Check:
 Checking firewall access: PASSED
 Validating certificate PASSED

Account Validation:
 Creating temporary profile: PASSED
 Creating temporary bucket: PASSED

S3 API Validation:
 Validating Put Bucket: PASSED
 Validating List Bucket: PASSED
 Validating Put Object: PASSED
 Validating Get Object: PASSED
 Validating List Object: PASSED
 Validating Delete Object: PASSED
 Validating Bulk Delete: PASSED

Cleaning Up:
 Deleting temporary bucket: PASSED
 Deleting temporary profile: PASSED

Provider verification passed.
```

12. Configure the file migration policy for this MTree. You can specify multiple MTrees in this command. The policy can be based on the age threshold or the range.

a. To configure the age-threshold (migrating files older than the specified age to cloud):

```
# data-movement policy set age-threshold age_in_days to-tier cloud cloud-unit
unitname mtrees mtreename
```

b. To configure the age-range (migrating only those files that are in the specified age-range):

```
# data-movement policy set age-range min-age age_in_days max-age age_in_days to-tier
cloud cloud-unit unitname mtrees mtreename
```

13. Export the file system, and from the client, mount the file system and ingest data into the active tier. Change the modification date on the ingested files such that they now qualify for data migration. (Set the date to older than the age-threshold value specified when configuring the data-movement policy.)
14. Initiate file migration of the aged files. Again, you can specify multiple MTrees with this command.

```
# data-movement start mtrees mtreename
```

To check the status of data-movement:

```
# data-movement status
Data-movement to cloud tier:
----------------------------
Data-movement was started on Aug 29 2020 07:45 and completed on Aug 31 2020 11:31
Moved (post-comp): 39.72 GiB, (pre-comp): 38.93 GiB,
Files inspected: 39881, Files eligible: 39881, Files moved: 39868, Files failed: 13

Data-movement recall:
---------------------
No recall operations found.
Use 'path all' option to list all recall jobs (up to the maximum supported limit).


To check the detailed status of data-movement:

# data-movement status to-tier cloud detailed
Data-movement to cloud tier:
----------------------------
Data-movement:
   90% complete; Elapsed time:  0:04:42
Moved (post-comp): 5.54 MiB, (pre-comp): 30.00 GiB,
Files inspected: 30, Files eligible: 30, Files moved: 0, Files failed: 0

Data-movement status for MTrees:
--------------------------------
MTree                     Files      Files   Files    Files   Bytes Moved   Bytes Moved
Destination
                      Inspected   Eligible   Moved   Failed    (Pre-comp)   (Post-comp)
Cloud Unit
------------------   ---------   --------   -----   ------   -----------   -----------
-----------
/data/col1/stu-av1*         10         10       0        0     10.03 GiB      1.31 MiB
cu_ecs
/data/col1/stu-av2*         10         10       0        0     10.03 GiB      2.42 MiB
cu_ecs
/data/col1/stu-av3*         10         10       0        0     10.03 GiB      1.81 MiB
cu_ecs
------------------   ---------   --------   -----   ------   -----------   -----------
-----------
(*) Data-movement is in progress for the marked MTrees.

Files currently being moved:
----------------------------
Path Name                                    File Size       Logical          Logical
Destination   Elapsed Time
                                                          Bytes Moved   Bytes Verified
Cloud Unit        hh:mm:ss
-----------------------------------------   ----------   -----------   --------------
-----------   ------------
/data/col1/stu-av1/FILE-115113.0001.0008   228.68 MiB    228.68 MiB         2.03 MiB
cu_ecs            00:04:38
/data/col1/stu-av1/FILE-115114.0002.0008   261.23 MiB    261.23 MiB         2.06 MiB
cu_ecs            00:04:38
/data/col1/stu-av1/FILE-115113.0007.0008   229.16 MiB    229.16 MiB         6.01 MiB
cu_ecs            00:04:38
/data/col1/stu-av1/FILE-115113.0004.0008   261.96 MiB    261.96 MiB         2.11 MiB
cu_ecs            00:04:38
-----------------------------------------   ----------   -----------   --------------
-----------   ------------
```

You can also watch the progress of data-movement:

```
# data-movement watch
```

15. Verify that file migration worked and the files are now in the cloud tier:

```
# filesys report generate file-location path all
```

16. Once you have migrated a file to the cloud tier, you cannot directly read from the file (attempting to do so results in an error). The file can only be recalled back to the active tier. To recall a file to the active tier:

```
# data-movement recall path pathname
```

Configuring encryption for DD cloud units

Encryption can be enabled at three levels: System, Active Tier, and cloud unit. Encryption of the Active Tier is only applicable if encryption is enabled for the system. Cloud units have separate controls for enabling encryption.

Steps

1. Select **Data Management** > **File System** > **DD Encryption**.

   NOTE: If no encryption license is present on the system, the Add Licenses page is displayed.

2. In the DD Encryption panel, do one of the following:
   - To enable encryption for **Cloud Unit *x***, click **Enable**.
   - To disable encryption for **Cloud Unit *x***, click **Disable**.

   NOTE: You are prompted to enter security officer credentials to enable encryption.

3. Enter the security officer **Username** and **Password**. Optionally, check **Restart file system now**.
4. Click **Enable** or **Disable**, as appropriate.
5. In the File System Lock panel, lock or unlock the file system.
6. In the Key Management panel, click **Configure**.
7. In the Change Key Manager dialog, configure security officer credentials and the key manager.

   Cloud encryption is supported with either the Embedded Key Manager or an External key manager.
8. Click **OK**.
9. Use the DD Encryption Keys panel to configure encryption keys.

Information needed in the event of system loss

Once Cloud Tier is configured, record the following information about the system and store it in a safe location apart from the system. This information will be needed to recover the Cloud Tier data in case the system is lost.

NOTE: This process is designed for emergency situations only and will involve significant time and effort from the Dell engineering staff.

- Serial number of the original system
- System passphrase of the original system
- DDOS version number of the original system
- Cloud Tier profile and configuration information

Using DD Replicator with Cloud Tier

Collection replication is not supported on systems with Cloud Tier enabled.

Managed file replication and MTree replication are supported on Cloud Tier enabled systems. One or both systems can have Cloud Tier enabled. If the source system is Cloud Tier enabled, data may need to be read from the cloud if the file was already migrated to the Cloud Tier. A replicated file is always placed first in the Active Tier on the destination system even when Cloud

Tier is enabled. A file can be recalled from the Cloud Tier back to the Active Tier on the source MTree only. Recall of a file on the destination MTree is not allowed.

Files on a replication source system that have not yet been replicated to the destination system are skipped during cloud data movement.

NOTE: Files in the Cloud Tier cannot be used as base files for virtual synthetic operations. The incremental forever or synthetic full backups need to ensure that the files remain in the Active Tier if they will be used in virtual synthesis of new backups.

Using DD Virtual Tape Library (VTL) with Cloud Tier

On systems configured with Cloud Tier and DD VTL, the cloud storage is supported for use as the VTL vault. DD VTL does not support the option to store the vault from an MTree replication destination on cloud storage.

To use DD VTL tape out to cloud, license and configure the cloud storage first, and then select it as the vault location for the VTL.

DD VTL tape out to cloud provides additional information about using VTL with Cloud Tier.

Displaying capacity consumption charts for Cloud Tier

Three charts are available for displaying Cloud Tier consumption statistics—Space Usage, Consumption, and Daily Written.

Steps

1. Select **Data Management** > **File System** > **Charts**.
2. For **Chart**, select one of the following:
   - Space Usage
   - Consumption
   - Daily Written
3. For **Scope**, select **Cloud Tier**.
   - The Space Usage Tab displays space usage over time, in MiB. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as pre-compression used (blue), post-compression used (red), and the compression factor (green).
   - The Consumption Tab displays the amount of post-compression storage used and the compression ratio over time, which enables you to analyze consumption trends. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as capacity (blue), post-compression used (red), compression factor (green), cleaning (orange) and data movement (violet).
   - The Daily Written Tab displays the amount of data written per day. You can select a duration (one week, one month, three months, one year, or All). The data is presented (color-coded) as pre-compression written (blue), post-compression used (red), and the total compression factor (green).

Cloud Tier logs

If Cloud Tier suffers a failure of any kind, in configuration or operation, the system automatically creates a folder with a timestamp that is associated with the time of the failure.

Mount the `/ddvar/log/debug` directory to access the logs.

NOTE: The output of the `log list` command does not list all the detailed log files that are created for the Cloud Tier failure.

Migrate an existing Cloud Tier system to a new system

It is possible to migrate the system data from and older system model configured with Cloud Tier storage to a newer system model to improve performance and use new features.

CAUTION: The following prerequisites apply to Cloud Tier migration:

- **A valid replication license is required.**
- **Disable all replication activity on the migration source system.**
- **Stop all MFR activity on the migration source system prior to running the `migration send` command. MFR activity is supported while the migration is in progress but file copies in progress at the time of the `migration send` command are terminated. Starting new file copy operations is not permitted after the `migration commit` command is run on the migration source.**
- **Stop all MFR activity on the migration source system prior to running the `migration commit` command.**
- **Verify the Cloud Tier feature is in the same state on both the migration source and migration destination systems. Cloud Tier should be enabled on both systems, or disabled on both systems.**
- **HA should be disabled on the source and the destination during migration. HA can only be enabled after migration is finished/committed on both source and destination.**

This migration process migrates the Active Tier storage, and the locally-stored Cloud Tier metadata from the existing system to a new system. During the Cloud Tier migration, the source system operates in a restricted mode where the Active Tier storage is available for backup operations, but I/O on the Cloud Tier storage is not permitted.

The Cloud Tier migration consists of the following steps:

1. Copy Active Tier data from the existing system to the new system.
2. Copy Cloud Tier metadata from the existing system to the new system.
3. Disconnect the cloud bucket from the existing system.
4. Connect the cloud bucket to the new system.
5. Commit the migration operation.

The following operations are not permitted while the migration is in progress:

- Sending Active Tier data to Cloud Tier storage.
- Recalling data from Cloud Tier storage.
- Cleaning the Cloud Tier storage.
- Restoring files directly or reading from the Cloud Tier storage.
- File system cleaning on the source system.
- System sanitization on the source system.
- Enabling or disabling file system encryption.
- Enabling, disabling, or setting the embedded key manager or an external key manager.
- Creating, destroying, deleting, or syncing keys from the embedded key manager or an external key manager.

Starting a Cloud Tier migration

The procedure to initiate the Cloud Tier migration is only available through the CLI. This procedure is supported for migrating physical DD systems to DDVE instances, or DDVE instances to physical DD systems.

Prerequisites

Review the following guidelines before beginning the migration:

- **Free space**: Verify the current cloud tier usage on the source system is at or below 90%. The system migration cannot start if a critical space alert present. Since ingest is allowed during a system migration, the overall space on the source ssytem Active Ter should never outgrow the space allocated on the destination.
- **System passhrase**: Verify the system passphrase is the same on both the source and destination systems to avoid reconfiguring the passphrase if an error is reported after the migration starts.

- **Network trust between the source and destination systems**: Verify network trust is properly set between the source and destination systems. Enable bi-directional access for SSH ports and check any firewall rules defined on the source system that might affect connectivity with the destination system.
- **Cloud unit secret keys**: Configure the cloud units on the destination system uses the same secret keys as the cloud units on the source system. The *Configuring cloud units* section provides additional information.
- **Home directories**: The migration only replicates data, home directories must be created manually on the destination system.

Complete the following tasks on the new system before beginning the Cloud Tier migration operation:

CAUTION: HA should be disabled on the source and the destination during migration. HA can only be be enabled after migration is finished/committed on both source and destination.

1. If the migration pair consists of two physical DD systems, verify both systems are running DDOS release 7.3.0.5 or higher.
2. If the migration pair consists of a physical DD system and a DDVE instance (on prem or DD3300), verify both systems are running a DDOS release in the DDOS 7.7 family or higher.
3. Verify both the source and destination systems are running the same DDOS release.
4. Add a Cloud Tier license on the new system.
5. Add other feature licenses as required on the new system.
6. If a passphrase is configured on the existing system, set the same passphrase on the new system.

   NOTE: The passphrase store-on-disk setting should not be less secure on the destination than on the source.
7. If encryption is configured on the existing system, set the same encryption values including key manager settings and FIPS compliance on the new system.
8. If automatic key rotation is configured on the existing system, disable it before starting the migration. Reenable it on the new system after the migration.
9. If encryption is configured on the existing system, back up the key export files from the existing system.
10. If Retention Lock Compliance is enabled on the existing system, enable RLC on the new system.
11. Record the cloud profile and cloud unit information from the existing system.
12. Create the file system on the new system, but do not enable it.

Steps

1. Verify the file system is created by not enabled on the new system.

   ```
   # filesys status
   ```
2. Enable Cloud Tier on the new system.

   ```
   # cloud enable
   ```
3. Create the cloud profile on the new system with the same configuration as on the existing system.

   ```
   # cloud profile add <profile-name>
   ```
4. Create the cloud units on the new system with the same configuration as on the existing system.

   ```
   # cloud unit add <unit-name> profile <profile-name> for-migration
   ```
5. Stop all data movement operations on the existing system.

   ```
   # data-movement stop all
   ```
6. Stop all cloud cleaning and background deletion on the existing system.

   ```
   # cloud clean stop
   # cloud clean background-delete stop
   ```
7. Stop all garbage collection operations on the existing system.

   ```
   # filesys clean stop
   # system sanitize abort
   ```
8. Prepare the new system to be the Cloud Tier migration destination.

   ```
   # migration receive source-host <source-host-name> [keephostname]
   ```

   NOTE: If you want to rename the hostname of the destination system with the hostname of the source system, specify the optional **keephostname** parameter to update the destination system replication configurations with the source system hostname.
9. Start the Cloud Tier migration from the existing system.

   ```
   # migration send all destination-host <destination-host-name>
   ```

10. Monitor the migration status until it completes.

    # migration status

    NOTE: If a cloud unit becomes unavailable during the migration, the system reports that cloud unit as `unavailable`.

    # migration watch

11. Commit the migration on the new system.

    # migration commit

12. Commit the migration on the existing system.

    # migration commit

13. If the existing system contained DD Boost storage units, they will appear as regular MTrees on the new system. Convert them back to DD Boost storage units on the new system.

    # ddboost user assign *<ddboost-username>*

    # ddboost storage-unit modify *<storage-unit-name>* user *<ddboost-username>* skip-chown

    NOTE: Use `skip-chown` when the DD Boost user already has ownership of existing files and directories in the storage unit, or when transfer of ownership is not required. This option can provide a performance boost to the conversion operation.

14. Enable the file system on the new system.

    # filesys enable

15. Redirect backup applications to the new system.
16. Decommission the existing system.

    # filesys disable

    # cloud unit delete *<cloud-unit-name>*

    # filesys destroy

Update the replication configuration or decommission the old system

After the migration is complete, the old system can be decommissioned or continue in use. To remain in use, additional configuration changes are required on the new system.

About this task

Complete one of the following steps.

Steps

1. Modify the replication configuration as described in one of the following examples.
   If replication is configured between systems *A* and *B*, with *A* as the source and *B* as the destination, the following steps are required after migrating data from *A* to the new system *A1*:
   a. Modify the replication contexts on system *B* to specify system *A1* as the replication source.
   b. Restart the file system on system *A1* to automatically synchronize the replication contexts to system *B*.

      NOTE: Breaking and resynchronizing the replication contexts is not required.

   c. After synchronization is complete and both systems are up and running, optionally change the system *A1* hostname to the same value as system *A*.

   If replication is configured between systems *A* and *B*, with *A* as the source and *B* as the destination, the following steps are required after migrating data from *B* to the new system *B1*:
   a. Modify the replication contexts on system *A* to specify system *B1* as the replication destination.
   b. Break and resynchronize the replication contexts that were modified to resume synchronization between systems *A* and *B1*.

      NOTE: Breaking the replication contexts removes them from the system. Create the contexts again on the replication source before resynchronizing them.

   c. After synchronization is complete and both systems are up and running, optionally change the system *B1* hostname to the same value as system *B*.
2. If required, decommission the old system.

   # filesys disable

```
# cloud unit delete <cloud-unit-name>
# filesys destroy
```

Using the CLI to remove Cloud Tier

You can use the CLI to remove the Cloud Tier configuration.

Prerequisites

Delete all files in the cloud units before removing the Cloud Tier configuration from the system. Run the `filesys report generate file-location path all output-file file_loc` command to identify the files in the cloud units, and delete them from the NFS mount points of the MTrees.

NOTE: The command above creates the report `file_loc` in the `/ddr/var/` directory.

Steps

1. Disable the file system.

```
# filesys disable

This action will disable the file system.
Applications may experience interruptions
while the file system is disabled.
    Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please wait.............
The filesystem is now disabled.
```

2. List the cloud units on the system.

```
# cloud unit list
Name           Profile        Status
------------   ------------   ------
cloud_unit-1   cloudProfile   Active
cloud_unit-2   cloudProfile2  Active
------------   ------------   ------
```

3. Delete the cloud units individually.

```
# cloud unit del cloud_unit-1
This command requires authorization by a user having a 'security' role.
Please present credentials for such a user below.
        Username: secofficer
        Password:
Please enter sysadmin password:

This command irrevocably destroys all data
 in the cloud unit "cloud_unit-1".
        Are you sure? (yes|no) [no]: yes

ok, proceeding.

Destroying cloud unit "cloud_unit-1"
Cloud unit 'cloud_unit-1' deleted. The data in the cloud will be deleted asynchronously
on the filesystem startup.
```

4. Verify the delete operations are in progress.

```
# cloud unit list
Name           Profile        Status
------------   ------------   ------
cloud_unit-1   cloudProfile   Delete-Pending
```

```
cloud_unit-2   cloudProfile2  Delete-Pending
------------   ------------   ------
```

5. Restart the file system.

```
# filesys enable
Please wait.....................................
The filesystem is now enabled.
```

6. Run the `cloud unit list` command to verify that neither cloud unit appears.
   Contact Support if one or both cloud units still display with the status `Delete-Pending`.
7. Identify the disk enclosures that are assigned to Cloud Tier.

```
# storage show tier cloud

Cloud tier details:
Disk   Disks               Count   Disk     Additional
Group                              Size     Information
------ ------------------- ------  -------- ------------
dgX    2.1-2.15, 3.1-3.15  30      3.6 TiB
------ ------------------- ------  -------- ------------
Current cloud tier size: 0.0 TiB
Cloud tier maximum capacity: 108.0 TiB
```

8. Remove the disk enclosures from Cloud Tier.

```
# storage remove enclosures 2, 3

Removing enclosure 2...Enclosure 2 successfully removed.

Updating system information...done

Successfuly removed: 2 done

Removing enclosure 3...Enclosure 3 successfully removed.

Updating system information...done

Successfuly removed: 3 done
```

26

DD Retention Lock

This chapter includes:

Topics:

DD Retention Lock overview

When data is locked on an MTree that is enabled with DD Retention Lock, DD Retention Lock helps ensure that data integrity is maintained. Any data that is locked cannot be overwritten, modified, or deleted for a user-defined retention period of up to 70 years.

There are two DD Retention Lock editions:

- *DD Retention Lock Governance Edition* retains the functionality of DD Retention Lock prior to DDOS 5.2. You can use DD Retention Lock Governance to define retention policies on data that is to be retained for a specific period of time to meet internal IT governance policies implemented by the system administrator.
- *DD Retention Lock Compliance Edition*, when properly installed, configured, enabled, and administered, enables you to meet strict data permanence requirements of regulatory standards, such as those of SEC 17a-4(f).

  Files locked on a Data Domain or PowerProtect system using DD Retention Lock Compliance Edition software cannot be altered or destroyed before the retention period expires. DD Retention Lock Compliance Edition requires a security officer for implementation of policies. An audit log file is accessible by the administrator or security officer.

Each edition requires a separate, add-on license, and either or both can be used on a single system.

The retention-locking protocol is the same for both the DD Retention Lock Governance and Compliance Editions. The differences in use stem from the system behavior for the DD Retention Lock Compliance Edition, since it places strict restrictions to meet compliance requirements. The KB article *Data Domain: Retention Lock Frequently Asked Questions*, avilable from https://www.dell.com/support, provides additional information.

The DD Retention Lock Governance Edition does not require a security officer and provides a higher degree of flexibility for archive data retention.

For archive compliance storage requirements, retention-locked files can be replicated using DD Replicator to another Data Domain or PowerProtect system. If a retention-locked file is replicated, it remains retention locked on the destination system, with the same level of protection as the source file.

The *DD Boost for Open Storage Administration Guide* and *DD Boost for Partner Integration Administration Guide* provide additional information about Veritas NetBackup (NBU) Integrated Retention Lock for DD Boost.

The topics that follow provide additional information on DD Retention Lock.

Related concepts

System behavior with DD Retention Lock

DD Retention Lock protocol

Only files that are explicitly committed to be retention-locked files are retention locked on the protection system. Files are committed to be retention-locked files through client-side file commands issued while DD Retention Lock Governance or Compliance is enabled on the MTree containing the files.

NOTE: Linux, Unix, and Windows client environments are supported.

Files that are written to shares or exports that are not committed to be retained (even if DD Retention Lock Governance or Compliance is enabled on the MTree containing the files) can be modified or deleted at any time.

Retention locking prevents any modification or deletion of files under retention from occurring directly from CIFS shares or NFS exports during the retention period specified by a client-side *atime* update command. Some archive applications and backup applications can issue this command when appropriately configured. Applications or utilities that do not issue this command cannot lock files using DD Retention Lock.

Retention-locked files are always protected from modification and premature deletion, even if retention locking is subsequently disabled or if the retention-lock license is no longer valid.

You cannot rename or delete non-empty folders or directories within an MTree that is retention-lock enabled. However, you can rename or delete empty folders, directories, or MTrees, and create new ones.

The retention period of a retention-locked file can be extended (but not reduced) by updating the file's *atime*.

For both DD Retention Lock Governance and Compliance, once the retention period for a file expires, the file can be deleted using a client-side command, script, or application. However, the file cannot be modified even after the retention period for the file expires. The system never automatically deletes a file when its retention period expires.

Related concepts

Client-Side Retention Lock file control
Setting Retention Locking on a file
Extending Retention Locking on a file
Supported data access protocols

DD Retention Lock flow

The general flow of activities with DD Retention Lock.

1. Enable MTrees for DD Retention Lock Governance or Compliance retention locking using the DD System Manager or DDOS commands issued from the system console.
2. Commit files to be retention locked on the protection system using client-side commands issued by an appropriately configured archiving or backup application, manually, or via scripts.

   NOTE: Windows clients may need to download utility programs for DDOS compatibility.
3. Optionally, extend file retention times using client-side commands.
4. Optionally, delete files with expired retention periods using client-side commands.

Related concepts

Extending Retention Locking on a file
Client-Side Retention Lock file control
Deleting or expiring a file

Automatic retention lock

The automatic retention lock functionality allows you to set automatic values for the retention period, and the lock delay (the time before a file becomes locked) on a per MTree basis. The automatic retention lock settings apply to new files created on the MTree after the retention lock settings are configured. Existing files are not impacted. This feature is supported for both Retention Lock Compliance and Retention Lock Governance.

Set the automatic retention period to ensure that every new file created on the MTree will be automatically locked and retained for the specified amount of time.

Set the automatic lock delay on the MTree to allow a period of time where a new file can be modified before it gets locked.

Automatic retention lock is subject to the following limitations:

- Retention Lock must be re-applied manually to any files reverted when automatic retention lock is in use.
- MTree replication of an MTree with automatic retention lock enabled to a system with an earlier version of DDOS that does not support automatic retention lock, results in the locked files replicating to the target system as regular files.
- In Automatic Retention Lock, for the files which are being ingested, the `mtree retention-lock report generate` command may incorrectly report those files as locked as well report an incorrect cooling off period.
- After configuration, automatic retention lock cannot be disabled on Retention Lock Compliance or Retention Lock Governance-enabled systems.
- After it is set, the automatic retention period cannot be reset on Retention Lock Compliance or Retention Lock Governance-enabled systems.

Supported data access protocols

DD Retention Lock is compatible with industry-standard, NAS-based Write-Once-Read-Many (WORM) protocols, and integration is qualified with archive applications such as Symantec Enterprise Vault, SourceOne, Cloud Tiering Appliance, or DiskXtender. Customers using backup applications such as CommVault can also develop custom scripts to use DD Retention Lock.

The protocol support of DD Retention Lock is as follows:

- NFS is supported with both DD Retention Lock Governance and Compliance.
- CIFS is supported with both DD Retention Lock Governance and Compliance.
- Automatic retention lock is supported on NFS and CIFS with both Retention Lock Governance and Compliance.
- DD VTL is supported with DD Retention Lock Governance, but not with DD Retention Lock Compliance. Automatic retention lock is not supported on DD VTL.

  Virtual tapes, here referred to as *tapes*, are represented as files on the file system.

  - You can retention-lock one or more tapes using the `vtl tape modify` command, described in the *DDOS Command Reference Guide*.

    The `mtree retention-lock revert` *path* command can be used to revert the retention-locked state of tapes that are locked with the `vtl tape modify` command. After the tape is unlocked, updates can be made to it. The unlocked state will not be visible via the DD System Manager or CLI until the DD VTL service is disabled then enabled. However, updates are applied to the unlocked tape. This capability is only for the DD Retention Lock Governance Edition.

  - The retention time for tapes can be displayed using the `vtl tape show` command with the `time-display retention` argument.
  - You can retention-lock an individual tape using the DD System Manager.
- vDisk (ProtectPoint Block Services) is supported with DD Retention Lock Governance but not with DD Retention Lock Compliance. Automatic Retention Lock is not supported on vDisk.

- DD Boost is supported with both DD Retention Lock Governance and Compliance. Automatic retention lock is supported on DD Boost with both Retention Lock Governance and Retention Lock Compliance. The *DD Boost for Partner Integration Administration Guide* provides additional information.

  If client-side scripts are used to retention-lock backup files or backup images, and if a backup application (Veritas NetBackup, for example) is also used on the system via DD Boost, be aware that the backup application may not share the context of the client-side scripts. Thus, when a backup application attempts to expire or delete files that were retention locked via the client-side scripts, space is not released on the Data Domain or PowerProtect system.

  Dell recommends that administrators change their retention period policy to align with the retention lock time. This applies to many of the backup applications that are integrated with DD Boost, including Veritas NetBackup, Veritas Backup Exec, and NetWorker.

  Setting retention lock during data ingest to a DD BOOST file in DSP mode is not allowed, and the client setting the RL receives an error. Retention lock should be set after the data ingest is complete.

  Setting retention lock during data ingest to a DD BOOST file in OST mode, or to an NFS file is not allowed and the client writing the data receives error as soon as RL is set. The partial file written before RL is set and committed to disk as a worm file.

Related concepts

DD Virtual Tape Library overview

Compliance mode on iDRAC

DD3300, DD6400, DD6900, DD9400, and DD9900 systems require that compliance mode be enabled on iDRAC before Retention Lock Compliance can be configured on the system.

Navigate to **Administration** > **Compliance** to view, and enable or disable iDRAC compliance access for DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

The **iDRAC Users** table displays the iDRAC users currently configured on the system, the role for each user, whether access for that user is enabled or disabled, and the amount of time those users will be allowed to access the system..

Create an iDRAC user account

Enable compliance mode on iDRAC for DD3300, DD6400, DD6900, DD9400, and DD9900 systems to use DD Retention Lock Compliance.

Prerequisites

This task is only for DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

Configure a security officer authorization policy on the system, and run the `system retention-lock configure` command to configure Retention Lock Compliance Edition on the system.

Steps

1. Select **Administration** > **Compliance**.

   NOTE: If no license for DD Retention Lock Compliance is present on the system, the **Compliance page displays with a message warning that there is no license.**

2. Click **Enable Retention Lock Compliance**.

   NOTE: This button is only available if Retention Lock Compliance Edition has been configured.

3. Specify the security officer credentials, and click **Enable**.
4. Create one or more iDRAC user accounts.
   a. In the **Role** list box, select **Operator (Disabled)** or **Read Only (Enabled)**.

      NOTE: Creating an iDRAC user with administrative privileges requires physical access to the system.

   b. In the **Username** field, specify a username for the iDRAC user account.

c. In the **Password** and **Confirm Password** fields, specify a password for the iDRAC user account.

NOTE: The password must comply to the iDRAC password policy, which may be different than the protection system password policy. The *Integrated Dell Remote Access Controller 9 User Guide* provides additional information about the iDRAC password policy.

d. Click **Add User** to add the user.
e. Specify details for another user account, or click **Save** to proceed.

Request PowerProtect access for iDRAC operators

Request PowerProtect access for iDRAC operator users for DD3300, DD6400, DD6900, DD9400, and DD9900 systems to use DD Retention Lock Compliance.

Prerequisites

This task is only for DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

Steps

1. Select **Administration** > **Compliance**.
2. Select an iDRAC operator from the **iDRAC Users** table.
3. Click **Enable**.
4. Specify the security officer credentials, and click **OK**.
5. In the **Duration** list box, select the amount of time to allow access and click **OK**.

NOTE: Unless the access duration is extended, the account will be automatically disabled once the specified duration expires.

Extend PowerProtect access for iDRAC operators

Extend PowerProtect access for iDRAC operator users for DD3300, DD6400, DD6900, DD9400, and DD9900 systems when access is required for a longer period of time.

Prerequisites

This task is only for DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

Steps

1. Select **Administration** > **Compliance**.
2. Select an iDRAC operator from the **iDRAC Users** table.
3. Click **Enable** and select the duration from the list box.
4. Specify the security officer credentials, and click **Authorize**.
5. In the **Duration** list box, select the amount of time to allow access and click **Save**.
6. Click **Yes** at the confirmation prompt.

Disable PowerProtect access for iDRAC operators

Disable PowerProtect access for iDRAC users for DD3300, DD6400, DD6900, DD9400, and DD9900 systems when access is no longer required.

Prerequisites

This task is only for DD3300, DD6400, DD6900, DD9400, and DD9900 systems.

About this task

When an iDRAC user is disabled or deleted, active iDRAC sessions for that user are automatically terminated.

Steps

1. Select **Administration** > **Compliance**.
2. Select an iDRAC operator from the **iDRAC Users** table.
3. Click **Disable**.
4. Specify the security officer credentials, and click **OK**.

Enabling DD Retention Lock on an MTree

Only files within DD Retention Lock Governance or Compliance enabled MTrees can be retention-locked.

MTrees enabled for DD Retention Lock Compliance cannot be converted to DD Retention Lock Governance MTrees and vice versa.

DD Retention Lock Governance is supported on:

- Physical DD systems (DDOS)
- DD3300 systems (DDVE)
- On-premises DDVE instances
- Cloud-based DDVE instances

DD Retention Lock Compliance is supported on:

- Physical DD systems (DDOS)
- DD3300 systems (DDVE)
- On-premises DDVE instances
- Cloud-based DDVE instances

The procedures that follow show how to enable MTrees for either DD Retention Lock Governance or DD Retention Lock Compliance.

Enabling DD Retention Lock Governance on an MTree

Add a DD Retention Lock Governance license to a system, and then enable DD Retention Lock Governance on one or more MTrees.

Steps

1. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.
   a. Select **Administration** > **Licenses**
   b. In the Licenses area click **Add Licenses**.
   c. In the License Key text box, type the license key.

      NOTE: License keys are case-insensitive. Include the hyphens when typing keys.

   d. Click **Add**.
2. Select an MTree for retention locking.
   a. Select **Data Management** > **MTree**.
   b. Select the MTree you want to use for retention locking. You can also create an empty MTree and add files to it later.
3. Click the MTree Summary tab to display information for the selected MTree.
4. Scroll down to Retention Lock area and click **Edit** to the right of Retention Lock.
5. Enable DD Retention Lock Governance on the MTree and change the default minimum and maximum retention lock periods for the MTree, if required.

   Perform the following actions in the Modify Retention Lock dialog box:

   a. Select **Enabled** to enable DD Retention Lock Governance on the MTree.
   b. In the **Use** drop-down list, select **Manual** or **Automatic**.

- For manual retention lock, to change the minimum or maximum retention period for the MTree:
  i. Type a number for the interval in the text box (for example, **5** or **14**).
  ii. From the drop-down list, select an interval (minutes, hours, days, years).
  NOTE: Specifying a minimum retention period of less than 12 hours, or a maximum retention period longer than 70 years, results in an error.
- For automatic retention lock, to change the minimum, maximum, or automatic retention period, or the automatic lock delay for the MTree:
  i. Type a number for the interval in the text box (for example, **5** or **14**).
  ii. From the drop-down list, select an interval (minutes, hours, days, years).
  NOTE: Specifying a minimum retention period of less than 12 hours, a maximum retention period longer than 70 years, an automatic retention period that does not fall between the minimum and maximum values, or an automatic lock delay less than 5 minutes or more than 7 days results in an error.

  NOTE: If a file is modified before the automatic lock delay has elapsed, the lock delay time starts over when the file modification is complete. For example, if the lock delay is 120 minutes and the file is modified after 60 minutes, the lock delay will start again at 120 minutes after the file is modified.

c. Click **OK** to save the settings.

After you close the Modify Retention Lock dialog box, which is updated MTree information appears in the Retention Lock area.

6. Check retention lock information for the MTree.

Note the following retention lock fields:

- Top:
  - The Status field indicates the read/write access for the MTree, the type of retention locking on the MTree, and whether retention locking is enabled or disabled.
- Bottom:
  - The Status field indicates whether retention locking is enabled for the MTree.
  - The Retention Period field indicates minimum and maximum retention periods for the MTree. The retention period that is specified for a file in the MTree must be equal to or greater than the minimum retention period and equal to or less than the maximum retention period.
  - The UUID field is a unique identification number that is generated for the MTree.

NOTE: To check retention lock configuration settings for any MTree, select the MTree in the Navigation Panel, then click the Summary tab.

Next steps

Retention-lock files in a retention-lock-enabled MTree.

Related concepts

Client-Side Retention Lock file control

Enabling DD Retention Lock Compliance on an MTree

Add a DD Retention Lock Compliance license to a system, set up a system administrator and one or more security officers, configure and enable the system to use DD Retention Lock Compliance software, and then enable DD Retention Lock Compliance on one or more MTrees.

About this task

Enabling Retention Lock Compliance on a DD6400, DD6900, DD9400, or DD9900 system locks down the iDRAC UI and SSH interfaces. Do not use the iDRAC interfaces to create additional iDRAC users because DDOS automatically disables those new users and reboots the system. After configuring Retention Lock Compliance, run the `user idrac create` command to create iDRAC users before running the `system retention-lock compliance enable` command.

Ensure that the system time is configured correctly or NTP is configured and synched before enabling DD Retention Lock Compliance. If the system limits on date change frequency and duration are set, they are enforced once DD Retention Lock

Compliance is enabled. Setting system date change frequency and date change limit provides additional information about the system limits on date change frequency and duration.

Steps

1. Add the DD Retention Lock Compliance license on the system, if it is not present.
   a. First, check whether the license is already installed.
      `elicense show`
   b. If the RETENTION-LOCK-COMPLIANCE feature is not displayed, install the license.
      `elicense update` *license-file*
2. Set up one or more security officer users accounts according to Role-Base Access Control (RBAC) rules.
   a. In the system administrator role, add a security officer account.
      `user add` *user* `role security`
   b. Enable the security officer authorization.
      `authorization policy set security-officer enabled`
3. Configure and enable the system to use DD Retention Lock Compliance.

   NOTE: Enabling DD Retention Lock Compliance enforces many restrictions on low-level access to system functions used during troubleshooting. Once enabled, the only way to disable DD Retention Lock Compliance is to initialize and reload the system, which results in destroying all data on the system.

   a. Configure the system to use DD Retention Lock Compliance.
      `system retention-lock compliance configure`
      The system automatically reboots.
   b. After the restart process is complete, create iDRAC users.
      `user idrac create`
   c. Enable DD Retention Lock Compliance on the system.
      `system retention-lock compliance enable`
4. Enable compliance on an MTree that will contain retention-locked files.
   `mtree retention-lock enable mode compliance mtree` *mtree-path*

   NOTE: Compliance cannot be enabled on /backup or pool MTrees.

5. To change the default minimum and maximum retention lock periods for a compliance-enabled MTree, type the following commands with security officer authorization.
   - `mtree retention-lock set min-retention-period` *period* `mtree` *mtree-path*
   - `mtree retention-lock set max-retention-period` *period* `mtree` *mtree-path*

   NOTE: The retention *period* is specified in the format [number] [unit]. For example: 1 min, 1 hr, 1 day, 1 mo, or 1 year. Specifying a minimum retention period of less than 12 hours, or a maximum retention period longer than 70 years, results in an error.

6. To change the automatic retention period and automatic lock delay for a compliance-enabled MTree, type the following commands with security officer authorization.
   - `mtree retention-lock set automatic-retention-period` *period* `mtree` *mtree-path*

     NOTE: The automatic retention *period* is specified in the format [number] [unit]. For example: 1 min, 1 hr, 1 day, 1 mo, or 1 year. The value must be between the minimum and maximum retention periods.
   - `mtree retention-lock set automatic-lock-delay` *time* `mtree` *mtree-path*

     NOTE: The automatic lock delay *time* is specified in the format [number] [unit]. For example: 5 min, 2 hr, or 1 day. The value must be between five minutes and seven days. The default is 120 minutes. If a file is modified before the automatic lock delay has elapsed, the lock delay time starts over when the file modification is complete. For example, if the lock delay is 120 minutes and the file is modified after 60 minutes, the lock delay will start again at 120 minutes after the file is modified.

   Repeat steps 4 through 6 to enable additional MTrees.

Next steps

Retention lock files reside in a retention-lock-enabled MTree.

Related concepts

Client-Side Retention Lock file control
Role-based access control

Place Indefinite Retention Hold (IRH) on an MTree

Place an IRH on an MTree to immediately restrict any modifications to data on that MTree for an indefinite period.

Prerequisites

DD Retention Lock (Governance or Compliance) must be enabled on the MTree to place an IRH.

About this task

IRH is not available for the default MTree, `/data/col1/backup`.

Steps

1. Select an MTree to place the IRH.
   a. Select **Data Management** > **MTree**.
   b. Select the MTree on which to place the IRH.
2. Click the MTree Summary tab to display information for the selected MTree.
3. Scroll down to the Indefinite Retention Hold area and click **Place Indefinite Retention Hold (IRH)**.
4. Click **OK**.

Results

The IRH:

- Prevents deletion or modification of files that are currently locked, or whose locks have expired.
- Prevents the disabling of DD Retention Lock Governance on an MTree with an IRH in place, while still allowing modification of Retention Lock attributes

Files that were never locked are not affected by the IRH.

Remove an IRH from an MTree

Remove an IRH from an MTree to allow deletion of expired files or disabling of DD Retention Lock Governance on that MTree.

About this task

Steps

1. Select an MTree from which to remove the IRH.
   a. Select **Data Management** > **MTree**.
   b. Select the MTree from which to remove the IRH.
2. Click the MTree Summary tab to display information for the selected MTree.
3. Scroll down to Indefinite Retention Hold area and click **Remove Indefinite Retention Hold**.
4. Click **OK**.

Results

File deletion is allowed immediately, and DD Retention Lock can be disabled on an MTree with the IRH removed.

Client-Side Retention Lock file control

This section describes the DD Retention Lock client command interface for locking files stored on the protection system. Client commands are the same for DD Retention Lock Governance and Compliance. Linux, Unix, and Windows client environments are supported; however, Windows clients may need to download utility programs with commands to lock files.

NOTE: If your application already supports industry-standard WORM, writing a WORM file to a DD Retention Lock Governance or Compliance enabled MTree will lock the file on the system. The retention time in the application should agree with the DD Retention Lock settings. You do not need to use the commands described in this section. To check whether an application is tested and certified for the DD Retention Lock, refer to the *DD Series with Archiving Software* compatibility matrix in the E-Lab Interoperability Navigator at https://elabnavigator.dell.com/eln/elnhome.

NOTE: Some client machines using NFS, but running a legacy OS, cannot set retention time later than 2038. The NFS protocol doesn't impose the 2038 limit and allows to specifying times until 2106. Further, DDOS doesn't impose the 2038 limit.

Client-side commands are used to manage the retention locking of individual files. These commands apply to all retention-lock-capable systems and must be issued in addition to the setup and configuration of DD Retention Lock on the system.

Required Tools for Windows Clients

You need the `touch.exe` command to perform retention-locking from a Windows-based client.

To obtain this command, download and install utilities for Linux/Unix-based applications from Microsoft according to your Windows version. These utilities are best recommendations from Dell and should be used per customer environment.

NOTE: The `touch` command for Windows may have a different format than the Linux examples in this chapter.

Follow the installation instructions provided and set the search path as needed on the client machine.

Client Access to System Files

After an MTree is enabled for DD Retention Lock Governance or Compliance, you can:

- Create a CIFS share based on the MTree. This CIFS share can be used on a client machine.
- Create an NFS mount for the MTree and access its files from the NFS mount point on a client machine.

NOTE: The commands listed in this section are to be used only on the client. They cannot be issued through the DD System Manager or CLI. Command syntax may vary slightly, depending on the utility you are using.

The topics that follow describe how to manage client-side retention lock file control.

Related concepts

Enabling DD Retention Lock on an MTree

Setting Retention Locking on a file

To perform retention locking on a file, change the last access time (*atime*) of the file to the desired retention time of the file, that is, the time when the file can be deleted.

This action is usually performed using the archive application, and all the archive applications that are qualified on the protection system today (per the *Data Domain Archive Application Compatibility Guide*) follow the basic locking protocol outlined here.

The future *atime* you specify must respect the minimum and maximum retention periods of the file's MTree (as offsets from the current time), as shown in the next figure.

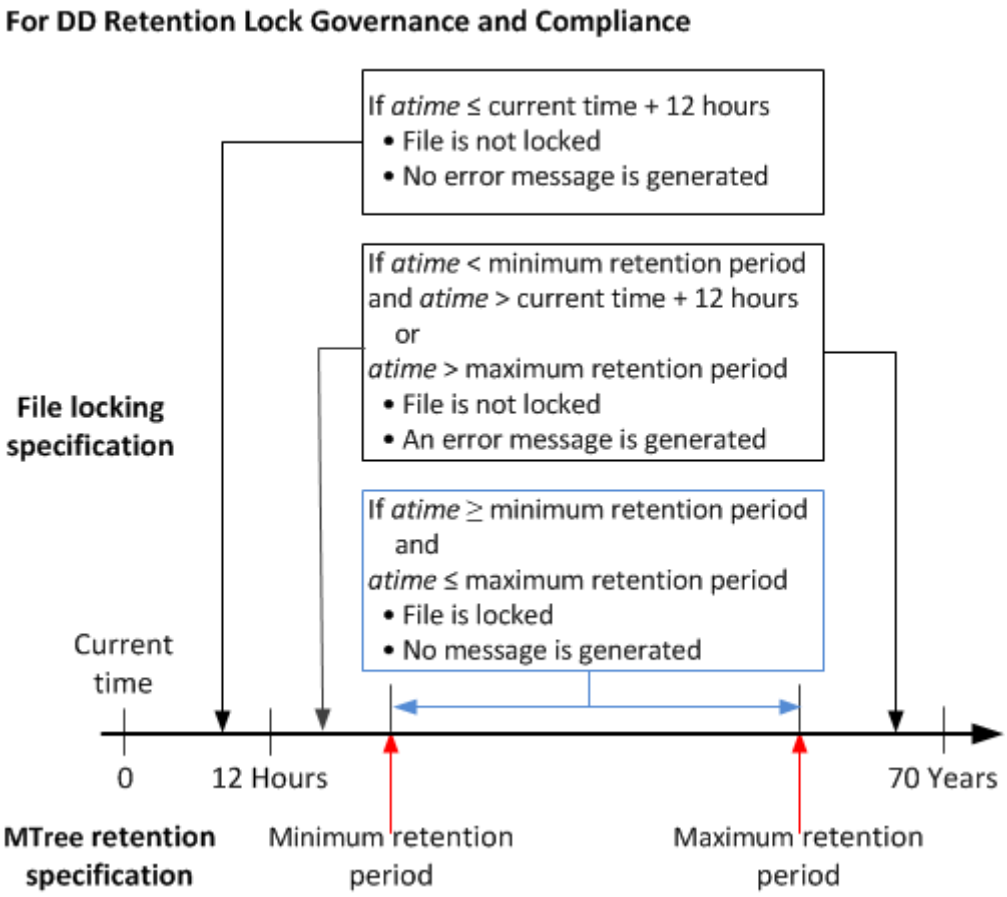

Figure 20. Valid and invalid *atimes* for retention locking files

NOTE: If a new retention lock time is smaller than the sum of the current retention lock time and the minimum retention period by three seconds or less, the system treats the new time as equal to the sum of the current time and the minimum retention period because the difference can occur due to latency.

NOTE: Some client machines using NFS, but running a legacy OS, cannot set retention time later than 2038. The NFS protocol does not impose the 2038 limit and allows to specifying times until 2106. Further, DDOS does not impose the 2038 limit.

Errors are permission-denied errors (referred to as EACCESS, a standard POSIX error). These are returned to the script or archive application setting the *atime*.

NOTE: A file must be completely written to the system before it is committed to be a retention-locked file.

The following command can be used on clients to set the *atime*:

```
touch -a -t [atime] [filename]
```

The format of *atime* is:

```
[[YY]YY] MMDDhhmm[.ss]
```

For example, suppose the current date and time is 1 p.m. on January 18, 2012 (that is, 201201181300), and the minimum retention period is 12 hours. Adding the minimum retention period of 12 hours to that date and time results in a value of 201201190100. Therefore, if the *atime* for a file is set to a value greater than 201201190100, that file becomes retention locked.

The following command:

```
ClientOS# touch -a -t 201412312230 SavedData.dat
```

will lock file `SavedData.dat` until 10:30 p.m. December 31, 2014.

Extending Retention Locking on a file

To extend the retention time of a retention-locked file, set the file's *atime* to a value greater than the file's current *atime* but less than the maximum retention period of the file's MTree (as an offset from the current time), as shown in the next figure.

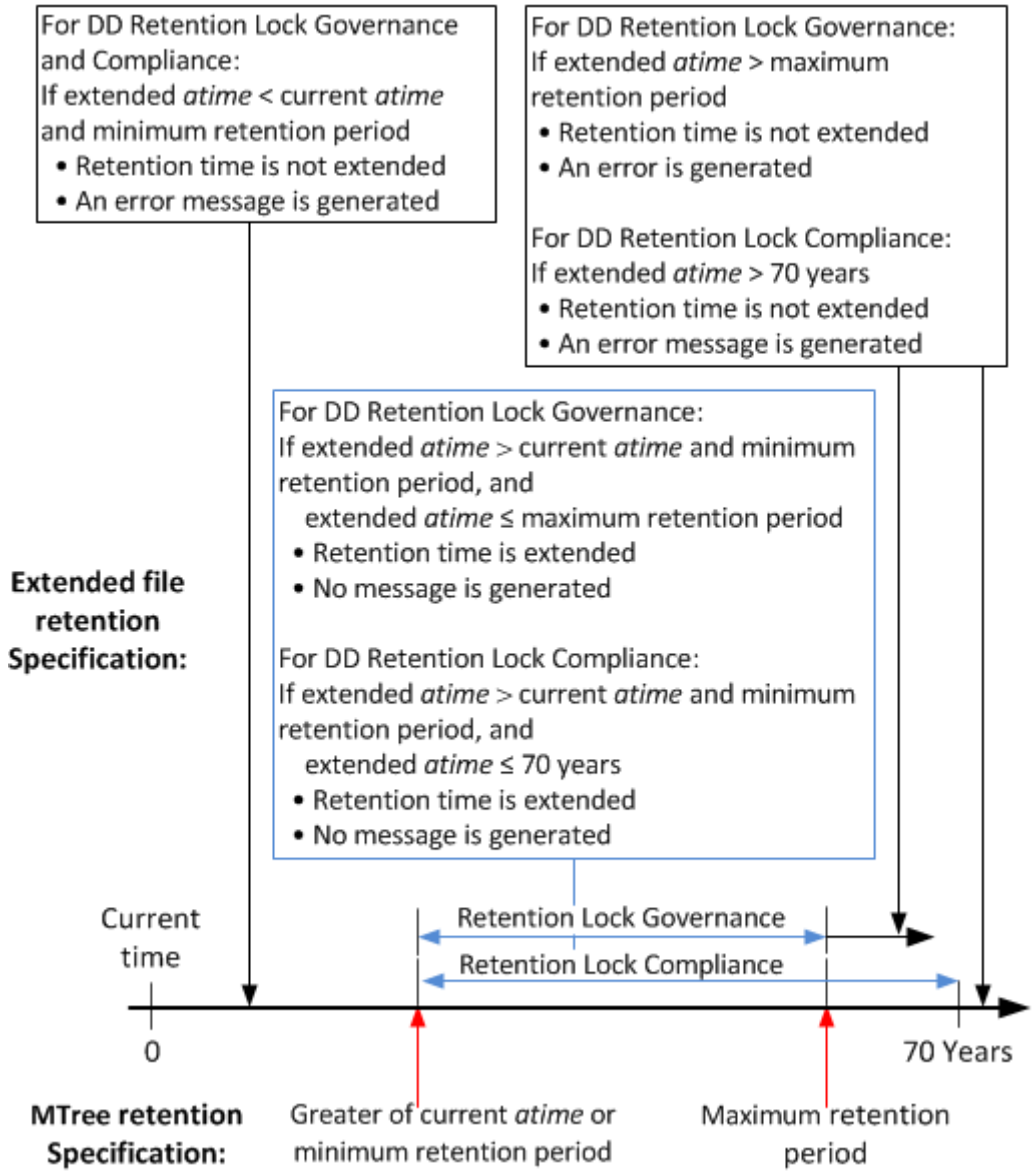

Figure 21. Valid and invalid *atimes* for extending retention locking on files

For example, changing the *atime* from `201412312230` to `202012121230` using the following command:

```
ClientOS# touch -a -t 202012121230 SavedData.dat
```

will cause the file to be locked until 12:30 p.m. December 12, 2020.

NOTE: Some client machines using NFS, but running a very old OS, cannot set retention time later than 2038. The NFS protocol doesn't impose the 2038 limit and allows to specifying times until 2106. Further, DDOS doesn't impose the 2038 limit.

Errors are permission-denied errors (referred to as EACCESS, a standard POSIX error). These are returned to the script or archive application setting the *atime*.

Identifying a Retention-Locked file

The *atime* value for a retention-locked file is its retention time. To determine whether a file is retention locked, try to set the *atime* of the file to a value earlier than its current *atime*. This action will fail with a permission-denied error if and only if the file is a retention-locked file.

First, list the current *atime* value, and then execute the `touch` command with an earlier *atime* using these commands:

```
ls -l --time=atime [filename]
touch -a -t [atime] [filename]
```

The following example shows the command sequence:

```
ClientOS# ls -l --time=atime SavedData.dat
202012121230
ClientOS# touch -a -t 202012111230 SavedData.dat
```

If the *atime* of `SavedData.dat` is 202012121230 (12:30 p.m. December 12, 2020) and the `touch` command specifies an earlier *atime*, 202012111230 (12:30 p.m. December 11, 2020), the `touch` command fails, indicating that `SavedData.dat` is retention-locked.

NOTE: The `--time=atime` option is not supported in all versions of Unix.

Specifying a directory and touching only those files

Use the command line to create a root directory containing the files for which access times will change.

In this routine, *root directory to start from* contains the files on which you want to change access times using this client system command:

```
find [root directory to start from] -exec touch -a -t [expiration time] {} \;
```

For example:

```
ClientOS# find [/backup/data1/] -exec touch -a -t 202012121230 {} \;
```

Reading a list of files and touching only those files

In this routine, *name of file list* is the name of a text file that contains the names of the files on which you want to change access times. Each line contains the name of one file.

Here is the client system command syntax:

```
touch -a -t [expiration time] 'cat [name of file list]'
```

For example:

```
ClientOS# touch -a -t 202012121230 'cat /backup/data1/filelist.txt'
```

Deleting or expiring a file

Delete or expire a file with an expired retention lock using a client application, or delete a file using a standard file-delete command.

Expiring a file using an application makes the file inaccessible to the application. The file may or may not actually be removed from the protection system by the expiration operation. If it is not removed, the application often provides a separate delete operation. You must have the appropriate access rights to delete the file, independent of DD Retention Lock.

NOTE: If the retention period of the retention-locked file has not expired, the delete operation results in a permission-denied error.

NOTE: For more information, refer to the KB article *Data Domain: How to delete data outside the backup application*, available from the Online Support website.

Privileged delete

For DD Retention Lock Governance (only), you can delete retention locked files using this two step process.

Steps

1. Use the `mtree retention-lock revert` *path* command to revert the retention locked file.

   NOTE: Only the sysadmin user can run `mtree retention-lock revert` command, and the command only works if DD Retention Lock Governance is enabled on the specified MTree.

2. Delete the file on the client system using the `rm` *filename* command.

Using ctime or mtime on Retention-Locked files

ctime is the last-metadata-change time of a file.

ctime

*ctime*gets set to the current time when any of the follow events occur:

- A non-retention-locked file is retention locked.
- The retention time of a retention-locked file is extended.
- A retention-locked file is reverted.

NOTE: User access permissions for a retention-locked file are updated using the Linux command line tool `chmod`.

mtime

mtime is the last-modified time of a file. It changes only when the contents of the file change. So, the mtime of a retention-locked file cannot change.

System behavior with DD Retention Lock

System behavior topics are discussed separately for DD Retention Lock Governance and DD Retention Lock Compliance in the sections that follow.

DD Retention Lock governance

Certain DDOS commands behave differently when using DD Retention Lock Governance. The following sections describe the differences for each.

Replication

Collection replication and MTree replication replicate the locked or unlocked state of files.

Files that are governance retention locked on the source are governance retention locked on the destination and have the same level of protection. For replication, the source system must have a DD Retention Lock Governance license installed—a license is not required on the destination system.

Collection replication and MTree replication replicate the minimum and maximum retention periods configured on MTrees to the destination system.

The procedure for configuring and using collection or MTree replication is the same as for protection systems that do not have a DD Retention Lock Governance license.

Replication Resync

The `replication resync` *destination* command tries to bring the destination into sync with the source when the MTree replication context is broken between destination and source systems. This command cannot be used with collection replication. Note that:

- If files are migrated to the cloud tier before the context is broken, the MTree replication resync overwrites all the data on the destination, so you will need to migrate the files to the cloud tier again.
- With Mtree replication, resync will fail if the source MTree does not have retention lock enabled and the destination MTree has retention lock enabled.
- With Mtree replication, resync will fail if the source and destination MTrees are retention lock enabled but the propagate retention lock option is set to FALSE.

Related concepts

DD Replicator overview

Fastcopy

When the `filesys fastcopy [retention-lock] source` *src* `destination` *dest* command is run on a system with a DD Retention Lock Governance enabled MTree, the command preserves the retention lock attribute during the fastcopy operation.

NOTE: If the destination MTree is not retention lock enabled, the retention-lock file attribute is not preserved.

Filesys destroy

Effects of the `filesys destroy` command when it is run on a system with a DD Retention Lock Governance enabled MTree.

- All data is destroyed, including retention-locked data.
- All `filesys` options are returned to their defaults. This means that retention locking is disabled and the minimum and maximum retention periods are set back to their default values on the newly created file system.

NOTE: This command is not allowed if DD Retention Lock Compliance is enabled on the system.

MTree delete

When the `mtree delete` *mtree-path* command attempts to delete a DD Retention Lock Governance enabled (or previously enabled) MTree that currently contains data, the command returns an error.

NOTE: The behavior of `mtree delete` is a similar to a command to delete a directory—an MTree with retention lock enabled (or previously enabled) can be deleted only if the MTree is empty.

Security officer authorization for `mtree retention-lock revert`

For systems with DD Retention Lock Governance, DDOS provides the ability to enable or disable the requirement for security officer authorization when reverting Retention Lock Governance MTrees.

With this option enabled, security officer authorization is required in addition to the sysadmin password to revert locked files. If the option is disabled, only the sysadmin password is required.

The security officer authorization option is disabled by default. Enabling the option requires both security officer authorization and the sysadmin password.

Enabling or disabling security officer authorization for reverting MTrees

Enabling and disabling security officer authorization for reverting Retention Lock Governance MTrees is only available through the CLI.

Steps

1. Enable or disable security officer authorization for reverting Retention Lock Governance MTrees.

```
# system retention-lock governance security-auth {enable | disable}
```

Enable security officer authorization for reverting Retention Lock Governance MTrees:

```
# system retention-lock governance security-auth enable
Please enter sysadmin password to confirm 'system retention-lock governance security-
auth enable | disable':
This command requires authorization by a user having a 'security' role.
Please present credentials for such a user below.
        Username: sec
        Password:
        Do you want to enable Security Officer authorization for retention-lock
governance? (yes|no) [no]: yes
Security Officer authorization has been enabled for retention-lock governance.

## mtree retention-lock revert /data/col1/mtree1
The 'mtree retention-lock revert' command removes retention-lock on this path thereby
making it unprotected.
Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please enter sysadmin password to confirm 'mtree retention-lock revert':
This command requires authorization by a user having a 'security' role.
Please present credentials for such a user below.
        Username: sec
        Password:
```

Disable security officer authorization for reverting Retention Lock Governance MTrees:

```
# system retention-lock governance security-auth disable
Please enter sysadmin password to confirm 'system retention-lock governance security-
auth enable | disable':
This command requires authorization by a user having a 'security' role.
Please present credentials for such a user below.
        Username: sec
        Password:
Do you want to disable the Security Officer authorization for Retention-lock governance?
(yes|no) [no]: yes
Security Officer authorization has been disabled for retention-lock governance.

# # mtree retention-lock revert /data/col1/mtree1
The 'mtree retention-lock revert' command removes retention-lock on this path thereby
making it unprotected.
Are you sure? (yes|no) [no]: yes

ok, proceeding.

Please enter sysadmin password to confirm 'mtree retention-lock revert':
```

2. Verify the status of security officer authorization for reverting Retention Lock Governance MTrees.

```
#  system retention-lock governance security-auth status
```

DD Retention Lock compliance

Certain DDOS commands behave differently when using DD Retention Lock Compliance. The following sections describe the differences for each.

Replication

An MTree enabled with DD Retention Lock Compliance can be replicated via MTree and collection replication.

MTree and collection replication replicate the locked or unlocked state of files. Files that are compliance retention locked on the source are compliance retention locked on the destination and have the same level of protection. Minimum and maximum retention periods configured on MTrees are replicated to the destination system.

To perform collection replication, the same security officer user must be present on both the source and destination systems before starting replication to the destination system and afterward for the lifetime of the source/replica pair.

Replication Resync

The `replication resync` *destination* command can be used with MTree replication, but not with collection replication.

- If the destination MTree contains retention-locked files that do not exist on the source, then resync will fail.
- Both source and destination MTrees must be enabled for DD Retention Lock Compliance, or resync will fail.

Related concepts

DD Replicator overview

Replication procedures

The topics in this section describe MTree and collection replication procedures supported for DD Retention Lock Compliance.

NOTE: For full descriptions of the commands referenced in the following topics, see the *DDOS Command Reference Guide*.

Replicating an MTree: One-to-one topology

Replicate a DD Retention Lock Compliance enabled MTree from a source system to a destination system.

Prerequisites

Enable DD Retention Lock on an MTree and configure client-side retention lock file control before replication.

Steps

1. Until instructed otherwise, perform the following steps on the destination system only.
2. Add the DD Retention Lock Compliance license on the system, if it is not present.
   a. First, check whether the license is already installed.
      `elicense show`
   b. If the RETENTION-LOCK-COMPLIANCE feature is not displayed, install the license.
      `elicense update` *license-file*
3. Set up one or more security officer users accounts according to Role-Base Access Control (RBAC) rules.
   a. In the system administrator role, add a security officer account.
      `user add` *user* `role security`
   b. Enable the security officer authorization.
      `authorization policy set security-officer enabled`
4. Configure and enable the system to use DD Retention Lock Compliance.

   NOTE: Enabling DD Retention Lock Compliance enforces many restrictions on low-level access to system functions used during troubleshooting. Once enabled, the only way to disable DD Retention Lock Compliance is to initialize and reload the system, which results in destroying all data on the system.

a. Configure the system to use DD Retention Lock Compliance.
```
system retention-lock compliance configure
```
The system automatically reboots.
b. After the restart process is complete, create iDRAC users.
```
user idrac create
```
c. Enable DD Retention Lock Compliance on the system.
```
system retention-lock compliance enable
```
5. Create a replication context.
```
replication add source mtree://source-system-name/data/col1/mtree-name destination
mtree://destination-system-name/data/col1/mtree-name
```
6. Perform the following steps on the source system only.
7. Create a replication context.
```
replication add source mtree://source-system-name/data/col1/mtree-name destination
mtree://destination-system-name/data/col1/mtree-name
```
8. Initialize the replication context.
```
replication initialize mtree://destination-system-name/data/col1/mtree-name
```
9. Confirm that replication is complete.
```
replication status mtree://destination-system-name/data/col1/mtree-name detailed
```
This command reports 0 pre-compressed bytes remaining when replication is finished.

Related concepts

Client-Side Retention Lock file control

Related tasks

Enabling DD Retention Lock Compliance on an MTree

Replicating an MTree: One-to-many topology

Replicate a DD Retention Lock Compliance enabled MTree from a source system to multiple destination systems.

Prerequisites

Enable DD Retention Lock compliance on an MTree and configure client-side retention lock file control before replication.

Steps

1. Until instructed otherwise, perform the following steps on the destination system only.
2. Add the DD Retention Lock Compliance license on the system, if it is not present.
a. First, check whether the license is already installed.
```
elicense show
```
b. If the RETENTION-LOCK-COMPLIANCE feature is not displayed, install the license.
```
elicense update license-file
```
3. Set up one or more security officer users accounts according to Role-Base Access Control (RBAC) rules.
a. In the system administrator role, add a security officer account.
```
user add user role security
```
b. Enable the security officer authorization.
```
authorization policy set security-officer enabled
```
4. Configure and enable the system to use DD Retention Lock Compliance.

NOTE: Enabling DD Retention Lock Compliance enforces many restrictions on low-level access to system functions used during troubleshooting. Once enabled, the only way to disable DD Retention Lock Compliance is to initialize and reload the system, which results in destroying all data on the system.

a. Configure the system to use DD Retention Lock Compliance.
```
system retention-lock compliance configure
```
The system automatically reboots.
b. After the restart process is complete, create iDRAC users.

```
user idrac create
```

c. Enable DD Retention Lock Compliance on the system.

```
system retention-lock compliance enable
```

5. Create a replication context.
   `replication add source mtree://source-system-name/data/col1/mtree-name destination mtree://destination-system-name/data/col1/mtree-name`
6. Perform the following steps on the source system only.
7. Create a replication context for each destination system.
   `replication add source mtree://source-system-name/data/col1/mtree-name destination mtree://destination-system-name/data/col1/mtree-name`
8. Initialize the replication context for each destination system MTree.
   `replication initialize mtree://destination-system-name/data/col1/mtree-name`
9. Confirm that replication is complete for each destination system.
   `replication status mtree://destination-system-name/data/col1/mtree-name detailed`

   This command reports 0 pre-compressed bytes remaining when replication is finished.

Related concepts

Client-Side Retention Lock file control

Related tasks

Enabling DD Retention Lock Compliance on an MTree

Adding DD Retention Lock Compliance protection to an existing MTree replication pair

Add DD Retention Lock Compliance protection to an existing MTree replication pair that is not enabled for retention locking.

Steps

1. Until instructed otherwise, perform the following steps on both the source and destination systems.
2. Log in to the DD System Manager.
   The DD System Manager window appears with **DD Network** in the Navigation panel.
3. Select a protection system.
   In the Navigation panel, expand **DD Network** and select a system
4. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.
   a. Select **Administration** > **Licenses**
   b. In the Licenses area click **Add Licenses**.
   c. In the License Key text box, type the license key.

      NOTE: License keys are case-insensitive. Include the hyphens when typing keys.

   d. Click **Add**.
5. Break the current MTree context on the replication pair.
   `replication break mtree://destination-system-name/data/col1/mtree-name`
6. Create the new replication context.
   `replication add source mtree://source-system-name/data/col1/mtree-name destination mtree://destination-system-name/data/col1/mtree-name`
7. Perform the following steps on the source system only.
8. Select an MTree for retention locking.
   Click the **Data Management** > **MTree** tab, then the checkbox for the MTree you want to use for retention locking. (You can also create an empty MTree and add files to it later.)
9. Click the MTree Summary tab to display information for the selected MTree.
10. Lock files in the compliance-enabled MTree.
11. Check the progress of resync.
    `replication watch mtree://destination-system-name/data/col1/mtree-name`
12. Confirm that replication is complete.

`replication status mtree://destination-system-name/data/col1/mtree-name detailed`

This command reports 0 pre-compressed bytes remaining when replication is finished.

Related concepts

Client-Side Retention Lock file control

Related tasks

Enabling DD Retention Lock Governance on an MTree

Converting a collection replication pair to MTree replication pairs

A procedure for customers who used collection replication under DD Retention Lock Compliance in DDOS 5.2 and want to convert compliance-enabled MTrees in the collection replication pair to MTree replication pairs.

Steps

1. On the source system only:
   a. Create a snapshot for each DD Retention Lock Compliance enabled MTree.
   `snapshot create snapshot-name /data/col1/mtree-name`
   b. Synchronize the collection replication pair.
   `replication sync col://destination-system-name`
   c. Confirm that replication is complete.
   `replication status col://destination-system-name detailed`

   This command reports 0 pre-compressed bytes remaining when replication is finished.
   d. View snapshot information for each DD Retention Lock Compliance enabled MTree.
   `snapshot list mtree /data/col1/mtree-name`
   Note the snapshot names for use later.
2. On the destination system only:
   a. Confirm that the replication is complete.
   `replication status mtree://destination-system-name/data/col1/mtree-name detailed`
   This command reports 0 pre-compressed bytes remaining when replication is finished.
   b. View each MTree snapshot replicated to the destination system.
   `snapshot list mtree /data/col1/mtree-name`
   c. Ensure that all DD Retention Lock Compliance MTree snapshots have been replicated by comparing the snapshot names generated here with those generated on the source system.
   `snapshot list mtree /data/col1/mtree-name`
3. On the both the source and destinations systems:
   a. Disable the file system.
   `filesys disable`
   b. Break the collection replication context.
   `replication break col://destination-system-name`
   c. Enable the file system. (Security officer authorization may be required.)
   `filesys enable`
   d. Add a replication context for each DD Retention Lock Compliance enabled MTree.
   `replication add source mtree://source-system-name/data/col1/mtree-name destination mtree://destination-system-name/data/col1/mtree-name`

   NOTE: Source and destination MTree names must be the same.
4. On the source system only:
   a. Check the progress of resync.
   `replication watch destination`
   b. Confirm that replication is complete.
   `replication status mtree://destination-system-name/data/col1/mtree-name detailed`
   This command reports 0 pre-compressed bytes remaining when replication is finished.

Performing collection replication

Replicate /data/col1 from a compliance-enabled source system to a compliance-enabled destination system.

About this task

NOTE: For collection replication the same security officer account must be used on both the source and destination systems.

Steps

1. Until instructed to do differently, perform the following steps on the source system only.
2. Log in to the DD System Manager.
   The DD System Manager window appears with **DD Network** in the Navigation Panel.
3. Select a protection system.
   In the Navigation Panel, expand **DD Network** and select a system.
4. Add the DD Retention Lock Governance license, if it is not listed under Feature Licenses.
   a. Select **Administration** > **Licenses**
   b. In the Licenses area click **Add Licenses**.
   c. In the License Key text box, type the license key.
      NOTE: License keys are case-insensitive. Include the hyphens when typing keys.
   d. Click **Add**.
5. Create the replication context.
   `replication add source col://source-system-name destination col://destination-system-name`
6. Until instructed to do differently, perform the following steps on the destination system only.
7. Destroy the file system.
   `filesys destroy`
8. Log in to the DD System Manager.
   The DD System Manager window appears with **DD Network** in the Navigation Panel.
9. Select a protection system.
   In the Navigation Panel, expand **DD Network** and select a system.
10. Create a file system, but do not enable it.
    `filesys create`
11. Create the replication context.
    `replication add source col://source-system-name destination col://destination-system-name`
12. Configure and enable the system to use DD Retention Lock Compliance.
    a. `system retention-lock compliance configure`
    b. `user idrac create`
    c. `system retention-lock compliance enable`
13. Perform the following steps on the source system only.
14. Initialize the replication context.
    `replication initialize col://destination-system-name`
15. Confirm that replication is complete.
    `replication status col://destination-system-name detailed`
    This command reports 0 pre-compressed bytes remaining when replication is finished.

Related tasks

Enabling DD Retention Lock Governance on an MTree

Adding DD Retention Lock Compliance protection to an existing collection replication pair

Add DD Retention Lock Compliance protection to a collection replication pair that was created without DD Retention Lock Compliance enabled on the source and destination systems.

Steps

1. Until instructed otherwise, perform the following steps on both the source and destination systems.
2. Disable the replication.
   `replication disable col://destination-system-name`
3. Log in to the DD System Manager.
   The DD System Manager window appears with **DD Network** in the Navigation Panel.
4. Select a protection system.
   In the Navigation Panel, expand **DD Network** and select a system.
5. Until instructed otherwise, perform the following steps on the source system.
6. Configure and enable the system to use DD Retention Lock Compliance.
   `system retention-lock compliance configure`
   (The system automatically reboots by executing the `system retention-lock compliance enable` command.)
7. Enable the replication context.
   `replication enable col://destination-system-name`
8. Until instructed otherwise, perform the following steps on the destination system.
9. Configure and enable the system to use DD Retention Lock Compliance.
   a. `system retention-lock compliance configure`
   b. `user idrac create`
   c. `system retention-lock compliance enable`
10. Enable the replication context.
    `replication enable col://destination-system-name`

Related tasks

Enabling DD Retention Lock Governance on an MTree

Fastcopy

When the `filesys fastcopy [retention-lock] source` *src* `destination` *dest* command is run on a system with a DD Retention Lock Compliance enabled MTree, the command preserves the retention lock attribute during the fastcopy operation.

NOTE: If the destination MTree is not retention lock enabled, the retention-lock file attribute is not preserved.

CIFS

When DD Retention Lock Compliance is enabled, CIFS servers no longer synchronize the system time with Active Directory. If there is a time difference of greater than five minutes between the system and Active Directory, the CIFS server displays an error message when an Active Directory user attempts to log in, or the system attempts to join an Active Directory domain. Configure Active Directory time with NTP to avoid this error.

CLI usage

Considerations for a protection system with DD Retention Lock Compliance.

- Commands that break compliance cannot be run. The following commands are disallowed:
  - `cloud unit del <unit-name>`
  - `authorization policy set security-officer {enabled | disabled}`
  - `filesys destroy`
  - `user reset`

- The following commands are not allowed on Retention Lock Compliance-enabled MTrees, but are allowed on Governance-enabled MTrees when Retention Lock Compliance is enabled on the system:
  - `mtree retention-lock disable mtree mtree-path`
  - `mtree retention-lock revert mtree-path`
- The following command requires security officer authorization if the license being deleted is for DD Retention Lock Compliance:
  - `elicense reset`
  - `elicense update`
- The following commands require security officer authorization if DD Retention Lock Compliance is enabled on an MTree specified in the command:
  - `mtree delete mtree-path`
  - `mtree rename mtree-path new-mtree-path`
  - `mtree retention-lock reset {min-retention-period period | max-retention-period period} mtree mtree-path`
  - `mtree retention-lock set {min-retention-period period | max-retention-period period} mtree mtree-path`
- The following commands require security officer authorization if DD Retention Lock Compliance is enabled on the system:

  NOTE: These commands must be run in interactive mode.

  - `alerts notify-list reset`
  - `config set timezone zonename`
  - `config reset timezone`
  - `cifs set authentication active-directory realm { [dc1 [dc2 ...]]`
  - `ntp add timeserver time server list`
  - `ntp del timeserver time server list`
  - `ntp disable`
  - `ntp enable`
  - `ntp reset`
  - `ntp reset timeservers`
  - `replication break {destination | all}`
  - `replication disable {destination | all}`
  - `system set date MMDDhhmm[[CC]YY]`

System clock

DD Retention Lock Compliance implements an internal security clock to prevent malicious tampering with the system clock.

The security clock closely monitors and records the system clock. If the accumulated skew between the security clock and the system clock reaches a designated value, the file system is disabled and can be resumed only by a security officer.

The allowed skew value is user-configurable, with a system default of 14 days if no value is specified.

Finding the System Clock Skew

You can run the DDOS command `system retention-lock compliance status` (security officer authorization required) to get system and security clock information, including the last recorded security clock value, and the accumulated system clock variance. This value is updated every 10 minutes.

Configure the system clock skew threshold

Configure the allowed clock skew threshold before the file system is shutdown and locked.

Prerequisites

Changing the allowed clock skew threshold from the system default of 14 days requires that a value be set for the system date change limit. Setting system date change frequency and date change limit describes how to set the system date change limit.

About this task

When Retention lock compliance is enabled, the clock-violation action can be set only once during the lifetime of the system. It can be modified multiple times before RLC is enabled, but after RLC is enabled, only one modification to the threshold is permitted.

This task must be performed through the CLI.

Steps

1. Run the `system show clock-violation-action` command to display the allowed clock skew value.

   If this parameter has not been configured, the system default value is 14 days.
2. Run the `system set clock-violation-action` command to set the allowed clock skew value.

   ```
   #  system set clock-violation-action {filesys-disable <threshold>}
   ```

   Specify the threshold in the format **[<*number*>] [<*unit*>]**. Possible unit values are:
   - min
   - hr
   - day
   - mo
   - year

   The minimum allowed value is the system date change limit, and the maximum allowed value is one year. The value **never** can be used to ensure the file system is never disabled for any amount of clock skew on the system.

Removing the system clock skew

Clock skew is updated every time the security clock records a new value for the system clock.

About this task

At any time, you can run the DDOS command `system set date` *MMDDhhmm*[[*CC*]*YY*] to set the time of the system clock (security officer authorization required). If the clock skew becomes larger than the preset value (two weeks or the threshold configured with the `system set clock-violation-action` command), the file system is disabled. Complete these steps to restart the file system and remove the skew between security and system clocks.

Steps

1. At the system console, enable the file system.
   `filesys enable`
2. At the prompt, confirm that you want to quit the `filesys enable` command and check whether the system date is right.
3. Display the system date.
   `system show date`
4. If the system date is not correct, set the correct date (security officer authorization is required) and confirm it.
   `system set date` *MMDDhhmm*`[[`*CC*`]`*YY*`]system show date`
5. Enable the file system again.
   `filesys enable`
6. At the prompt, continue to the enabling procedure.
7. A security officer prompt appears. Complete the security officer authorization to start the file system. The security clock will automatically be updated to the current system date.

NTP

When DD Retention Lock Compliance is enabled, some NTP functionality is restricted to admin users and offsets between the system time and NTP time causes some `ntp` commands to fail and generate alerts.

The following `ntp` commands require admin privileges to run:

- `ntp add timeserver`
- `ntp del timeserver`

- `ntp disable`
- `ntp enable`
- `ntp reset`
- `ntp reset timeservers`
- `ntp sync`

The following actions are triggered by offsets between the system time and the NTP time:

- If the offset is greater than the minimum of the system date change limit or 1000 seconds:
  - The `ntp enable` command fails and generates an alert.
  - The `ntp add timeserver` and `ntp sync` commands fail and generate alerts.
- If the offset between DD system time and NTP timeservers configured in DHCP is greater than the minimum of the system date change limit or 1000 seconds, the `ntp reset` command fails and raises an alert.

The system administrator is not allowed to run the `ntp sync` command more often than the system date change frequency. Any attempt to do so fails and generates an alert.

All operations to add, delete, or reset NTP timeservers generate alerts.

Successful operations to disable and synchronize NTP generate alerts.

27

DD Encryption

This chapter includes:

Topics:

DD Encryption overview

Data encryption protects user data if the protection system is stolen or if the physical storage media is lost during transit, and it eliminates accidental exposure of a failed drive if it is replaced.

When data enters the protection system using any of the supported protocols (NFS, CIFS, DD VTL, DD Boost, and NDMP Tape Server), the stream is segmented, fingerprinted, and de-duplicated (global compression). It is then grouped into multi-segment compression regions, locally compressed, and encrypted before being stored to disk (or object storage for cloud-based DDVE instances).

Once enabled, the DD Encryption feature encrypts all data entering the system. You cannot enable encryption at a more granular level.

CAUTION: Data that has been stored before the DD Encryption feature is enabled does not automatically get encrypted. To protect all of the data on the system, be sure to enable the option to encrypt existing data when you configure encryption.

Additional Notes:

The `filesys encryption apply-changes` command applies any encryption configuration changes to all data present in the file system during the next cleaning cycle. For more information about this command, see the *DDOS Command Reference Guide*.

DD Encryption supports all of the currently supported backup applications described in the Backup Compatibility Guides available through Online Support at https://www.dell.com/support.

DD Replicator can be used with encryption, enabling encrypted data to be replicated using collection, MTree, or application-specific managed file replication with the various topologies. Each replication form works uniquely with encryption and offers the same level of security. For more information, see the section on using DD Encryption with replication.

Files locked using DD Retention Lock can be stored, encrypted, and replicated.

The autosupport feature includes information about the state of encryption on the system:

- Whether or not encryption is enabled
- The Key Manager in effect and which keys are used
- The encryption algorithm that is configured
- The state of the file system

Related concepts

Using DD Encryption with DD Replicator

Configuring encryption

This procedure includes configuring a key manager.

If the Encryption Status on the **Data Management** > **File System** > **Encryption** tab shows Not Configured, click **Configure** to set up encryption on the protection system.

NOTE: The system passphrase must be set to enable encryption.

Provide the following information:

- Algorithm
  - Select an encryption algorithm from the drop-down list or accept the default AES 256-bit (CBC).

    The AES 256-bit Galois/Counter Mode (GCM) is the most secure algorithm but it is slower than the Cipher Block Chaining (CBC) mode.
  - Determine what data is to be encrypted: existing and new or only new. Existing data will be encrypted during the first cleaning cycle after the file system is restarted. Encryption of existing data can take longer than a standard file system cleaning operation.
- Key Manager (select one of the two)
  - Embedded Key Manager

    By default, the protection system Embedded Key Manager is in effect after you restart the file system.

    You can enable or disable key rotation. If enabled, type a rotation interval between 1 to 12 months.
  - External key manager

  NOTE: See the section about key management for an explanation about how the Embedded Key Manager and external key managers work.

The Summary shows the selected configuration values. Review them for correctness. To change a value, click **Back** to browse to the page where it was entered and modify it.

NOTE: Applications may experience an interruption while the file system is restarted.

Related concepts

Managing the system passphrase

Related tasks

Changing key managers after setup

About key management

Encryption keys determine the output of the cryptographic algorithm. They are protected by a passphrase, which encrypts the encryption key before it is stored in multiple locations on disk. The passphrase is generated by the user and requires both an administrator and a security officer to change it.

A key manager controls the generation, distribution, and life cycle management of multiple encryption keys. A protection system can use either the Embedded Key Manager, or an external key manager.

Only one can be in effect at a time. When encryption is enabled on a protection system, the Embedded Key Manager is in effect by default. If you configure the external key manager, it replaces the Embedded Key Manager and remains in effect until you disable it.

The Embedded Key Manager provides and generates multiple keys internally, although the system uses only one key at a time to encrypt data coming into the system.

The system supports a maximum of 1024 encryption keys, and allows you to specify a key rotation policy to set how many weeks or months a key is in effect before it is replaced. Manage key manager key rotation from the protection system.

KMIP-compliant external key managers

DD appliances support a KMIP-compliant key manager: KeySecure v8.5, v8.9, v8.10 and v8.12.1; NextGen v1.9.1 and v.1.10 from SafeNet/Gemalto, Data Security Manager (DSM) 6.3 from Vormetric/Thales, CipherTrust 2.1, 2.2, 2.3, and 2.4 from Thales, and Security Guardium Key Lifecycle Manager (GKLM) 4.1.1.0 from IBM. To use a KMIP key manager, users have to configure both the key manager and the protection system/DDVE, to trust each other. A protection system will retrieve these keys and their states from the key manager after establishing a secure TLS connection.

Rectifying lost or corrupted keys

Create a file that contains all of your system's current encryption keys. Your support provider can use this file to import keys back to your system should they become lost or corrupted. It is recommended that you create an export file on a regular basis.

You are prompted for the Security Officer's credential to export the keys. For additional key file protection, you can use a passphrase that differs from the one used in a protection system. After exporting, it is recommended that you save the key file in a secure file server accessible only by authorized users. You must remember the passphrase used for the key file. If the passphrase is lost or forgotten, the protection system cannot import and restore the keys. Enter:

```
# filesys encryption keys export
```

Key manager support

All Key Managers support all DDOS file system protocols.

Replication

When configuring protection systems for MTree replication, configure each system separately. The two systems can use either the same or a different key class, and the same or different key managers.

For collection replication configuration, the protection system must be configured on the source. All replicated data is encrypted with the key set on the source. New data written to the destination after a replication break will either use the last active key set on the source, or a new key if the key manager is configured.

Working with the Embedded Key Manager

When the Embedded Key Manager is selected, the protection system creates its own keys.

After the key rotation policy is configured, a new key is automatically created at the next rotation. You can disable the key rotation policy by clicking the disable button that is associated with the Embedded Key Manager Key's rotation status.

Creating a key (Embedded Key Manager)

Create an encryption key for the Embedded Key Manager.

Steps

1. Select **Data Management** > **File System** > **DD Encryption**.
2. In the Encryption Keys section, click **Create...** .
3. Type your security officer user name and password.
   A new protection system key is created and activated immediately.
4. Click **Create**.

Destroying a key (Embedded Key Manager)

Destroy an encryption key for the Embedded Key Manager.

Steps

1. Select **Data Management** > **File System** > **Encryption**.
2. In the Encryption Keys section, select the key in the list to be destroyed.
3. Click **Destroy...**.
   The system displays the Destroy dialog that includes the tier and state for the key.
4. Type your security officer user name and password.
5. Confirm that you want to destroy the key by clicking **Destroy**.

   NOTE: After a file system clean has run, the key state changes to Destroyed.

Deleting a key

You can delete Key Manager keys that are in the Destroyed or Compromised-Destroyed states. However, you only need to delete a key when the number of keys has reached the maximum 1024 limit. This procedure requires security officer credentials.

About this task

NOTE: To reach the Destroyed state, the Destroying a Key procedure must be performed on the key and a system cleaning must be run.

Steps

1. Select **Data Management** > **File System** > **Encryption**.
2. In the Encryption Keys section, select the key or keys in the list to be deleted.
3. Click **Delete...**.
   The system displays the key to be deleted, and the tier and state for the key.
4. Type your security officer user name and password.
5. Confirm that you want to delete the key or keys by clicking **Delete**.

Working with an external key manager

The DD system supports external key managers by using the Key Management Interoperability Protocol (KMIP) and centrally manages encryption keys in a single, centralized platform.

- When applicable, keys will be pre-created on the key manager.
- External key management is supported on both Active Tier and Cloud Tier storage.

Key manager setup

Follow the instructions for the type of key manager you are using.

Setting up KMIP key manager

With KMIP support, a protection appliance can retrieve symmetric key objects that are used for data at rest encryption from KMIP key managers.

Steps

1. Set up a KeySecure, DSM, CipherTrust, or GKLM instance with IP address <IP1>.
2. Create and install an SSL server certificate on the key manager.
3. Create a certificate signing request (CSR) for the system on the protection system/DDVE or Linux computer.
   a. Log in to the protection system.

b. Issue the command `adminaccess certificate cert-signing-request generate`.
If the command is successful, it generates the file `CertificateSigningRequest.csr`, which is located in `/ddvar/certificates/`.

By default, NFS exports do not have permissions to access the certificates folder, even to a root user.

```
# mount 16tbddve:/ddvar /mnt/DDVE
# cd /mnt/DDVE/certificates/
bash: cd: /mnt/DDVE/certificates/: Permission denied
# ls -al /mnt/DDVE/
total 800292
drwxr-xr-x 25 root staff      4096 Apr 10 08:32 .
drwxr-xr-x 26 root root       4096 Oct 24 12:11 ..
-rwxr-xr-x  1 root staff       180 Apr 10 08:36 .bashrc
drwxrwsr-x  2 root staff      4096 Aug 18  2016 benchmark
drwxr-sr-x  3 root staff      4096 Apr  4 15:49 cacerts
drwxrwsr-x  2 root staff      4096 Apr  4 12:50 cdes
drwxrws---  2 root staff      4096 Apr 11  2017 certificates
drwxrwsr-x  3 root staff      4096 Jul  1  2016 core
```

4. Take this CSR and have it issued/signed by the CA on the key manager.
If the command is successful, it generates the file `CertificateSigningRequest.csr`, which is located in `/ddvar/certificates/`.
5. Download that signed certificate (x.509 pem file] on to the protection system and use the private key of the CSR to create a pkcs#12 file.

   Rename`csr` to `pem` in the file name.
6. Download the root CA certificate from the CA that signed the key manager's KMIP server certificate.
7. On the protection system or DDVE, use `adminaccess` CLI to install the pkcs#12 client certificate and the CA certificate. Use application type as **keysecure**, **dsm**, **ciphertrust**, or **gklm**.
8. On the external key manager, create a symmetric key with AES-256 as the algorithm and key length.
   - For KeySecure KeyManager:
     a. Set the owner to the user that will use as KMIP on the protection system or DDVE.
     b. Select the `Exportable` option.
     c. Under **Security** > **Keys** > **Attributes** for the key, ensure to set **Application Namespace** to **DD_DARE_KEYS**. Ensure to set **Application Data** to key-class that you are planning to use on the protection system or DDVE.
   - For DSM:
     a. No action is required at this time. The `filesystem encryption key-manager set` command later in this task will request that DSM create a new key if it does not already have one.
   - For CipherTrust:
     a. Specify the KMIP user as the owner.
     b. Select the**Exportable** and **Deletable** options.
     c. Select **Keys** > **KMIP** > ***<Key-name>*** > **Properties** to verify the Application namespace is set as **`DD_DARE_KEYS`**.
     d. Verify that **Application Data** is set to the key-class used on the protection system or DDVE instance.
   - For GKLM:
     a. No action is required at this time. The `filesystem encryption key-manager set` command later in this task will request that GKLM create a new key if it does not already have one.
9. Use `filesys encryption key-manager set` command to configure ALL the parameters to access the key manager.
10. Enable the external key-manager by using the command `filesys encryption key-manager enable`.

    NOTE: If Cloud Tier encryption is enabled, the command prompts to enable a weekly key rotation policy.
11. Enable encryption by using the commands `filesys encryption enable`.
12. Keys should be automatically retrieved from the key-manager and seen in the local key table.
Sample output of local key table for `filesys encryption keys show`:

```
# filesys encryption keys show
Active Tier:
Key Key    State        Size
Id  MUID               post-comp
--- ---- ------------ ---------
1   d05  Activated-RW 40.50 MiB
--- ---- ------------ ---------
* Post-comp size is based on last cleaning of Tue Nov 5 06:42:54 2019.
```

The current active key is used to encrypt any data being ingested.

13. Sync the key states.
    a. On the KeySecure web interface, create a new active key as previously described.
    b. On the KeySecure web interface, deactivate the old key by clicking the key and going under the **Life Cycle** tab. Click **Edit State**. Set the **Cryptographic State** to **Deactivated**. Click **Save**.
14. On the protection system, sync the local key table by running the `filesys encryption keys sync` command. Sample output of local key table for`filesys encryption keys show`:

```
# filesys encryption keys show
Active Tier:
Key Key    State         Size
Id  MUID                post-comp
--- ---- ------------ ---------
1   d05  Activated-RW 40.50 MiB
--- ---- ------------ ---------
* Post-comp size is based on last cleaning of Tue Nov 5 06:42:54 2019.
```

NOTE: Keys can be marked as versioned keys. When 2nd and 3rd versions of a specific key are generated, KMIP queries currently don't pick up these keys and may be an issue if that key is being used by a protection system or DD VE.

Results

After key manager configuration is complete, the system will generate an alert if it loses the connection to the external key manager.

Using DD System Manager to set up and manage the KMIP-compliant key manager

This section describes how to use DD System Manager to configure and manage the KeySecure Key Manager, Data Security Manager, CipherTrust Key Manager, or Security Guardium Key Lifecycle Manager.

Configuring the external Key Manager

Use DD System Manager to set the key rotation policy from the protection system.

Prerequisites

Confirm the desired Key rotation period (weeks or months), the Key rotation start date, and the Next key rotation date.

Steps

1. Select **Data Management** > **File System** > **DD Encryption**.
2. In the **Key Management** section, click **Configure**. The **Change Key Manager** dialog box opens.
3. Enter your security officer user name and password.
4. Select **KeySecure Key Manager**, **Data Security Manager**, **CipherTrust Key Manager**, or **GKLM Key Manager** from the **Type** drop down menu. The Key Manager information appears.
5. Set the key rotation policy:

   NOTE: The rotation policy is specified in weeks and months. The minimum key rotation policy increment is one week, and the maximum key rotation policy increment is 52 weeks (or 12 months).

   a. Enable the Key Rotation policy. Set the **Enable Key rotation policy** button to enable.
   b. Enter the appropriate dates in the Key rotation schedule field.
   c. Select the appropriate number of weeks or months from the **Weeks** or **Months** drop down menu.
   d. Click **OK**.

Results

The key rotation policy is set or changed.

Creating a KMIP key

Create a KMIP encryption key for KeySecure Key Manager, DSM, CipherTrust, or GKLM.

Steps

1. Scroll down to the **Key Manager Encryption Keys** table.
2. Click **Add** to create a new Key Manager encryption key.
   a. Enter the Security Officer username and password.
   b. Click **Create**.

   A new KIMP key is created and activated immediately.

Modifying the state of an existing KMIP key

Use DD System Manager to modify the state of an existing KIMP encryption key for KeySecure Key Manager, DSM, CipherTrust, or GKLM.

Prerequisites

Review the conditions for changing a key state:

- A key in an `Activated-RO` key requires no conditions. Deactivate at any time.

Steps

1. Select **Data Management** > **File System** > **DD Encryption**.
2. Scroll down to view the **Key Manager Encryption Keys** table.
3. Select the appropriate key from the **Key Manager Encryption Keys** table.
4. To deactivate a key:
   a. Click on any key that shows an `Activated` state.
   b. Enter the security officer username and password.
   c. Click **DEACTIVATE**.

Figure 22. Change KMIP key to a Deactivated state

Results

The state of an existing key is changed.

Using the DD CLI to manage the external key manager

This section describes how to use the CLI to manage KeySecure Key Manager, DSM, CipherTrust, or GKLM.

Set or reset a KMIP key rotation policy

Use the CLI to set the key rotation policy on the system to periodically rotate keys. Note that the rotation policy is specified in weeks and months. The minimum key rotation policy increment is one week, and the maximum key rotation policy increment is 52 weeks (or 12 months).

Prerequisites

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Steps

1. Log into the system using the security role:
   Username: `sec`
   Password: `<security officer password>`

2. Set a key rotation policy for the first time. In our example, we will set the rotation policy to **three weeks**:

```
# filesys encryption key-manager set key-rotation-policy
                  {every <n> {weeks | months} | none}
```

For example:

```
# filesys encryption key-manager set key-rotation-policy every 3 weeks
```

Output that is similar to the following appears:

```
Key-rotation-policy is set. Encryption key will be rotated every 3 weeks.
```

3. Subsequently, run this command if you choose to change the existing key rotation policy. In our example, we will change the rotation policy from **three weeks** to **four months**:

NOTE: Log into the Data Domain system using the security role (where Username is `sec`, and the password is the `<security officer password>`).

```
# filesys encryption key-manager reset [key-rotation-policy]
```

For example:

```
filesys encryption key-manager set key-rotation-policy every 4 months
```

Output that is similar to the following appears:

```
Key-rotation-policy is set. Encryption key will be rotated every 4 months.
```

4. Display the current key rotation policy, or verify that the policy is set correctly:

```
# filesys encryption key-manager show
```

Output that is similar to the following appears:

```
The current key-manager configuration is:
Key Manager:                              Enabled
Server Type:                              <KeySecure/DSM/CipherTrust/GKLM>
Server:                                   <IP address of KMIP server>
Port:                                     5696
Status:                                   Online
Key-class:                                <key-class>
KMIP-user:                                <KMIP username>
Key rotation period:                      2 months
Last key rotation date:                   03:14:17 03/19 2018
Next key rotation date:                   01:01:00 05/17 2018
```

Results

The key rotation policy is set or changed.

Create a new active KMIP key

Use the protection system CLI to create a new active key.

Prerequisites

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Steps

1. Log into the protection system using the security role:
   Username: `<security office user>`
   Password: `<security officer password>`
2. Create a new active key:

   ```
   # filesys encryption key-manager keys create
   ```

3. Output that is similar to the following appears:

   ```
   New encryption key was successfully created.
   ```

Results

A new active key is created.

Modify the state of an existing KMIP key

Use the protection system CLI to modify the state of an existing key to a deactivated state.

Prerequisites

Ensure that you have the appropriate user credentials. The security role is required to run these commands.

Steps

1. Log into the protection system using the security role:
   Username: `sec`
   Password: `<security officer password>`
2. Modify the state of an existing key:

   ```
   # filesys encryption key-manager keys modify muid <key-muid> state deactivated
   ```

   For example:

   ```
   # filesys encryption key-manager keys modify muid
   740D711374A8C964A62817B4AD193C8DC44374A6ED534C85642782014F2E9D41 state deactivated
   ```

3. Output that is similar to the following appears:

   ```
   Key state modified.
   ```

Results

The state of an existing key is modified.

Handling key compromise scenarios with the external key manager

Encryption affects the performance of cleaning operations when data encrypted with the Compromised or Marked-For-Destroyed keys is re-keyed using the Activated-RW key.

After reencryption, no data remains encrypted with the Compromised or Marked-For-Destroyed keys. Also, any data written by the reencryption operation is encrypted with the Activated-RW key.

Reencrypting data

After compromising a key used to encrypt data, use the CLI to reencrypt the data.

Steps

1. After a key is compromised, run the `filesys encryption keys sync` command to synchronize the encryption keys.
2. Run the `cloud re-encrypt start` command to initiate the reencryption operation.
   After reencryption is complete, the compromised key is in the Destroyed-Compromised state.
3. Optionally run the `cloud re-encrypt watch` command to monitor the progress of the reencryption operation as it runs.
4. Optionally run the `cloud re-encrypt status` command to display the status of the previous reencryption operation.

Results

At the end of the reencryption operation, no data remains encrypted with the Compromised or Marked-For-Destroyed keys. All data written by the reencryption operation is encrypted with the Activated-RW key.

Reencrypting Active Tier data

Steps

1. After a key is compromised, run the `filesys encryption keys sync` command to synchronize the encryption keys.
2. Run the `filesys clean start` command to initiate the reencryption operation.
   After reencryption is complete, the compromised key is in the Destroyed-Compromised state.
3. Optionally run the `filesys clean watch` command to monitor the progress of the reencryption operation as it runs.
4. Optionally run the `filesys clean status` command to display the status of the previous reencryption operation.

Reencrypting cloud data

Steps

1. After a key is compromised, run the `filesys encryption keys sync` command to synchronize the encryption keys.
2. Run the `cloud re-encrypt start` command to initiate the reencryption operation.
   After reencryption is complete, the compromised key is in the Destroyed-Compromised state.
3. Optionally run the `cloud re-encrypt watch` command to monitor the progress of the reencryption operation as it runs.
4. Optionally run the `cloud re-encrypt status` command to display the status of the previous reencryption operation.

Results

If a large amount of data requires reencryption, running this procedure multiple times may be required because there is a hard limit on the amount of data reencrypted at one time. The **`cloud re-encrypt status`** command displays how many containers have been processed and how many still require reencryption. At the end of the reencryption operations, no data remains encrypted with the Compromised or Marked-For-Destroyed keys. All data written by the reencryption operation is encrypted with the Activated-RW key.

Changing key managers after setup

Change the key manager configured on the system between the embedded key manager or a supported KMIP key manager.

Prerequisites

To manage certificates for a system, you must start DD System Manager on that system.

Steps

1. Select **Data Management** > **File System** > **Encryption**.
2. Under Key Management, click **Configure**.
3. Type your security officer username and password.
4. Select a new key manager type.
5. Specify the server name, key class, and user.
6. If required, change the port number.

   NOTE: The default port is 5696.

7. Select to enable or disable key rotation. If enabled, enter a rotation interval between 1-to-12 months. Click **OK**.
8. Click **Manage Certificates** to add certificates.

Migrating keys between key managers

DDOS provides the ability to migrate keys from one supported external key manager to another.

Prerequisites

Complete the following before migrating keys:

- Run the `adminaccess certificate import` command to import the certificate for the new key manager.
- Configure the system to use the new key manager.
- From the external key managers, verify the keys are migrated from the old key manager to the new before running the migration command on the protection system.

About this task

Run the following command to migrate the keys from one key manager server to another after switching key managers:

Steps

```
filesys encryption key-manager keys migrate source {keysecure | dsm | ciphertrust | gklm}
destination {keysecure | dsm | ciphertrust | gklm}
```

```
#  filesys encryption key-manager keys migrate source keysecure destination ciphertrust
```

Deleting certificates

Select a certificate with the correct fingerprint.

Steps

1. Select a certificate to delete.
2. Click **Delete**.
   The system displays a Delete Certificate dialog with the fingerprint of the certificate to be deleted.
3. Click **OK**.

Checking DD Encryption settings

Check the settings for the DD Encryption feature.

Click the **Data Management** > **File System** > **Encryption** tabs. The currently used Key Manager is shown as Enabled. For a description of the DD Encryption settings, see the section about the encryption view.

Enabling and disabling DD Encryption

After configuring DD Encryption, the status is enabled and the Disabled button is active. When DD Encryption is disabled, the Enabled button is active.

Enabling DD Encryption

Use the DD System Manager to enable the DD Encryption feature.

Steps

1. Using the DD System Manager, select the protection system you are working with in the Navigation panel.
2. In the Encryption view, click the **Enable** button.
3. Both of the following options are available:
   - Select **Apply to existing data** and click **OK**. Encryption of existing data will occur during the first cleaning cycle after the file system is restarted.
   - Select **Restart the file system now** and click **OK**. DD Encryption will be enabled after the file system is restarted.

Next steps

NOTE: If the **Apply to existing data** option is not selected, only new data is encrypted, while existing data already stored on the system remains unencrypted. Run the **`filesys encryption apply-changes`** command followed by the **filesys clean start** command to encrypt the pre-existing unencrypted data at any point.

Applications may experience an interruption while the file system is restarted.

Disabling DD Encryption

Use the DD System Manager to disable the DD Encryption feature.

Steps

1. Using the DD System Manager, select the protection system you are working with in the Navigation panel.
2. In the Encryption view, click the **Disable** button.

   The Disable Encryption dialog box is displayed.
3. In the Security Officer Credentials area, enter the user name and password of a security officer.
4. Select one of the following:
   - Select **Apply to existing data** and click **OK**. Decryption of existing data will occur during the first cleaning cycle after the file system is restarted.
   - Select **Restart the file system now** and click **OK**. DD Encryption will be disabled after the file system is restarted.

Next steps

NOTE: Applications may experience an interruption while the file system is restarted.

Locking and unlocking the file system

Use this procedure when an DD Encryption-enabled protection system (and its external storage devices) are being transported, or if you want to lock a disk that is being replaced. The procedure requires two accounts: Security Officer and System Administration roles.

Steps

1. Select **Data Management** > **File System** > **Encryption** .

   In the File System Lock area, the Status shows whether the file system is Locked or Unlocked.
2. Disable the file system by clicking **Disabled** in the File System status area.
3. Use the procedure to lock or unlock the file system.

Locking the file system

To lock the file system, DD Encryption must be enabled and the file system must be disabled.

Steps

1. Select **Data Management** > **File System** > **Encryption** and click **Lock File System**.
2. In the text fields of the Lock File System dialog box, provide:
   - The username and password of a Security Officer account (an authorized user in the Security User group on that protection system).
   - The RSA passcode for the Security Officer account, if multifactor authentication is enabled.
   - The current and a new passphrase.
3. Click **OK**.

   This procedure re-encrypts the encryption keys with the new passphrase. This process destroys the cached copy of the current passphrase (both in-memory and on-disk).

   NOTE: Changing the passphrase requires two-user authentication to protect against the possibility of a rogue employee's shredding the data.

CAUTION: Be sure to take care of the passphrase. If the passphrase is lost, you will never be able to unlock the file system and access the data. The data will be irrevocably lost.

4. Shut down the system:

   CAUTION: Do not use the chassis power switch to power off the system. Type the following command at the command prompt instead.

   ```
   # system poweroff The 'system poweroff' command shuts down the system and turns off the power. Continue? (yes|no|?) [no]:
   ```

5. Transport the system or remove the disk being replaced.
6. Power on the system and use the procedure to unlock the file system.

Related tasks

Unlocking the file system

Unlocking the file system

This procedure prepares an encrypted file system for use after it has arrived at its destination.

Steps

1. Select **Data Management** > **File System** > **Encryption** and click **Unlock File System**.
2. In the text fields, type the passphrase that was used to lock the file system.
3. Click **OK**.
4. Click **Close** to exit.

   If the passphrase is incorrect, the file system does not start and the system reports the error. Type the correct passphrase, as directed in the previous step.

Related tasks

Locking the file system

Changing the encryption algorithm

Reset the encryption algorithm if necessary, or select options to encrypt new and existing data or just new data.

Steps

1. Select **Data Management** > **File System** > **Encryption**
2. To change the Encryption Algorithm used to encrypt the protection system, click **Change Algorithm**.

   The Change Algorithm dialog box is displayed. Supported encryption algorithms are:
   - AES-128 CBC
   - AES-256 CBC
   - AES-128 GCM
   - AES-256 GCM
3. Select an encryption algorithm from the drop-down list or accept the default AES 256-bit (CBC).

   The AES 256-bit Galois/Counter Mode (GCM) is the most secure algorithm but it is significantly slower than the Cipher Block Chaining (CBC) mode.

   NOTE: To reset the algorithm to the default AES 256-bit (CBC), click Reset to default.

4. Determine what data will be encrypted:
   - To encrypt existing and new data on the system, select **Apply to Existing data, Restart file system now**, and click **OK**.

     Existing data will be encrypted during the first cleaning cycle after the file system is restarted.

NOTE: Encryption of existing data can take longer than a standard file system clean operation.

- To encrypt only new data, select **Restart file system now** and click **OK**.

5. The status is displayed. Click **Close** when the process is complete.

NOTE: Applications may experience an interruption while the file system is restarted.