Dell EqualLogic Group Manager

Administrator’s Guide

PS Series Firmware Version 9.1

FS Series Firmware Version 4.0

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.

2017 - 03

Rev. 110-6269-EN-R1

Contents

About This Manual

Dell EqualLogic PS Series arrays optimize resources by automating capacity, performance, and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software and firmware updates. Dell EqualLogic FS Series appliances, when combined with PS Series arrays, offer a high-performance, high-availability, scalable NAS solution.

Audience

The PS Series and FS Series documentation is designed for administrators who are responsible for managing a PS Series group and one or more FS Series appliances. Administrators are not required to have extensive network or storage system experience. However, it is helpful to understand:

- Basic networking concepts
- Current network environment
- Application disk storage requirements

In addition, administrators might find it useful to understand the basic concepts for:

- Network topologies
- RAID configurations
- Disk storage management

Related Documentation

For detailed information about PS Series arrays, FS Series appliances, groups, volumes, array software, and host software, log in to the Documentation page at the Dell EqualLogic support site.

Dell Online Services

To learn about Dell products and services:

1. Visit dell.com or the URL specified in any Dell product information.
2. Use the locale menu, or click the link that specifies your country or region.

Dell EqualLogic Storage Solutions

To learn more about Dell EqualLogic products and new releases being planned, visit the following site: Dell EqualLogic TechCenter. At this site, you can also see articles, demos, online discussions, and more details about the features and benefits of the Dell product family.

Dell Technical Support and Customer Service

Dell support is available to answer your questions about PS Series arrays and FS Series appliances.

Contacting Dell

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services might not be available in your area. To contact Dell for sales, technical support, or customer service issues, go to dell.com/support.

About Group Manager

Group Manager is an easy-to-use SAN and NAS management tool integrated with the Dell EqualLogic PS Series firmware. Providing a comprehensive single point of management, Group Manager eliminates the need for a dedicated management workstation or server by enabling administrators to remotely manage virtually any aspect of their EqualLogic iSCSI-based SAN or NAS.

Group Manager provides detailed information about SAN and NAS configuration, and enables storage managers to perform group administration tasks such as provisioning, snapshots, replication scheduling, and other management requirements quickly and easily, using windows and dialog boxes. Built-in monitoring and notifications provide email, syslog support, comprehensive Simple Network Management Protocol (SNMP) monitoring and traps, and many more standard monitoring and notification features.

Figure 1. Group Manager GUI shows the Group Manager GUI interface for a typical group.

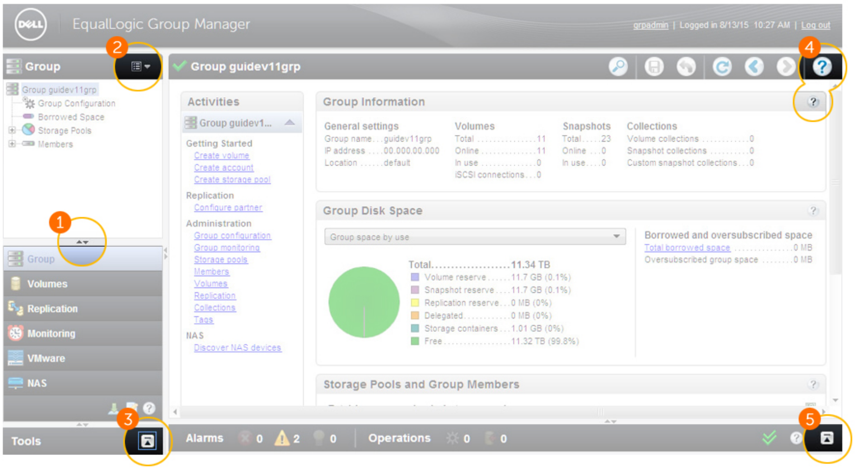

General settings
Group name . . . guidev11grp
IP address 00.000.00.000
Location . . . . . . default

Volumes
Total 11
Online 11
In use 0
iSCSI connections . . . 0

Snapshots
Total 23
Online . . . 0
In use 0

Collections
Volume collections 0
Snapshot collections 0
Custom snapshot collections . . . 0

Total. 11.34 TB
Volume reserve . . . . . . 11.7 GB (0.1%)
Snapshot reserve . . . . 11.7 GB (0.1%)
Replication reserve . . . 0 MB (0%)
Delegated . . . . 0 MB (0%)
Storage containers . . . 1.01 GB (0%)
Free . . . . 11.32 TB (99.8%)

Borrowed and oversubscribed space
Total borrowed space . . . . 0 MB
Oversubscribed group space . . . . 0 MB

Figure 1. Group Manager GUI

Callout	Function	Description
1	View Drag Handle	Drag this handle up or down to view options for accessing the areas you are working within. Expand the list completely to view the NAS option. Drag the handle down to collapse the list of selections at the bottom.
2	Tree View Options	Displays options that are specific to the section you are in (for example, volumes, groups).
3	Expand Tools View	Displays a list of tools that you can access.
4	Context Sensitive Help	Click on question marks to open the Group Manager's online help. The page where you click the question mark displays the help for that page.
5	Alarms and Operations Panel	Displays the Alarms and Operations panel, where you can monitor the status of and take action on various tasks that must be performed by the user. You can also view the reason for and recommended solution for the different alarms that might happen (for example, warning, critical).

About GUI and CLI Access

By default, PS Series group administrators can access the GUI remotely using a web browser or a standalone Java application. Administrators can also manage a group by using the CLI across a telnet, SSH or a serial connection. If you use a serial connection you must be connected to the primary controller.

The Group Manager graphical user interface (GUI) is based on the Java platform. You can access the interface from any computer with a web browser and a connection to the EqualLogic SAN. In addition to ports 3002 and 3003, the GUI uses the standard HTTP port (80).

Group Manager also provides a command-line interface (CLI), which allows you to perform the same operations as the Group Manager GUI. The Group Manager CLI provides a comprehensive set of commands for managing a PS Series group or FS Series NAS cluster. The CLI also enables you to manage individual group members for maintenance purposes.

You can disable network access to the GUI or CLI, preventing any administrator from logging in to the group or from using CLI commands.

If you disable all methods of access to the group, you must use a serial connection and the CLI to manage the group or to reenable access. See your *Hardware Owner's Manual* for information about serial connections.

Log In to the Group Manager GUI

You can access Group Manager's graphical user interface (GUI) from any web browser with access to the EqualLogic SAN.

1. In your web browser, type the group IP address.
2. Type the account name.
3. Type the password that was specified when the account was created.

You can connect to Group Manager using HTTP. The first time you log in to the Group Manager GUI using an account with group administrator privileges, the system prompts you to accept an end-user license agreement (EULA). The EULA also appears if any changes have been made to the agreement since the last time you logged in.

Access the Command-Line Interface (CLI)

To access the group to run CLI commands, use one of the following connections:

- Network connection

  From a computer, use telnet or SSH to connect to the group (or management) IP address or — if you are running array management commands on a specific array—connect to an IP address assigned to a network interface on the array.
- Serial connection

  Set up a serial connection to the array, as appropriate for the control module model. Make the connection to port 0 on the active control module (identified by a LED indicator labeled ACT that is illuminated green). Use the serial cable that shipped with the array. See the *Hardware Owner's Manual* for your array model for more information.

  You can also access the CLI on PS-M4110 systems with the Dell M1000e Chassis Management Console (CMC). For more information, see the hardware documentation included with your system.

For information about using the Group Manager CLI, refer to the *Dell EqualLogic Group Manager CLI Reference Guide*. This guide is located on the EqualLogic support website at eqlsupport.dell.com (login required).

About the Localized GUI

The Group Manager GUI is available in the following languages:

- French (fr)

- German (de)
- Japanese (ja)
- Korean (ko)
- Simplified Chinese (zh)
- Spanish (es)

The Group Manager GUI defaults to the same language as set for the browser and operating system. If your browser and operating system are set to a non-English supported language, and you want the GUI to display in English, log in to the Group Manager GUI using the Group Manager IP with **/english.html** at the end (for example, http://*ip_address*/**english.html**).

About the End-User License Agreement

The first time you log in (from either a browser or standalone application) to the Group Manager GUI using group administrator privileges, the system prompts you to accept an end-user license agreement (EULA). If you accept the EULA, Group Manager launches. If you do not accept the EULA, you are logged out of Group Manager. You need to accept the EULA only once per group. If you later add more members to the group, the system does not ask you again to accept the EULA.

If your configuration includes FS Series network-attached storage (NAS) appliances at the time that you accept the EULA, your acceptance of the EULA covers both your PS Series arrays and FS Series NAS appliances. If your configuration does not include FS Series NAS appliances when you accept the EULA, and you later attempt to create a NAS cluster by adding these appliances, the system prompts you to accept the EULA again to cover your the appliances. You need to accept the EULA only once for the FS Series NAS appliances. If you subsequently add more FS Series appliances to the NAS cluster, the system does not ask you again to accept the EULA.

Install the Group Manager GUI Locally

To control a specific PS Series group, you can install the Group Manager GUI on a computer and run it as a standalone application. The first time that you log in to the Group Manager GUI with group administrator privileges, the system prompts you to accept an end-user license agreement (EULA). You can install GUIs for more than one group on a single computer.

Prerequisite: Make sure the required Java version is installed on the computer on which you install the GUI application. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for details about Java versions and standalone GUI requirements.

When you install the GUI locally, it automatically updates when you update the PS Series firmware. However, you must log out of the GUI and then log in again after performing a firmware update to make sure the GUI displays all the features in the updated firmware.

If you change the IP address of a group for which you are running the GUI locally, or if you configure a management network for the group, you must uninstall the GUI application and then install it again.

To install the Group Manager GUI locally:

1. Click **Tools** in the lower-left portion of the GUI window and then click **Run as application**.
2. Confirm that you want to install the GUI application. Depending on the Java version, you might be prompted to create a shortcut on the desktop.

After installation, the standalone GUI starts automatically. You can log in to it as if you were accessing it with a web browser. For future GUI sessions, click or double-click the GUI icon to launch the GUI.

Show or Hide the Session Banner

You can customize the message on the sign-on banner page that displays when you log in to Group Manager.

NOTE: The custom banner page can contain up to 1000 *bytes* of text. The number of characters in the banner page varies with the number of bytes required for each character. Banner pages containing multibyte characters accommodate shorter messages.

To set the session banner:

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. Select the **Show banner before login** checkbox.
4. Click **Set session banner**.
5. Type the banner message in the field in the dialog box.

   You can also copy text and paste it into the dialog box.
6. Press **Return** to wrap the banner text.
7. (Optional) Click **Preview** to see how the banner will look.
8. Click **OK**.

To disable the session banner, clear the **Show banner before login** checkbox and save the change.

Keyboard Shortcuts

Table 1. Keyboard Shortcuts

Location	Action	Shortcut
General	Switch to Group view	Ctrl+Alt+G
	Switch to Volumes view	Ctrl+Alt+V
	Switch to Replication view	Ctrl+Alt+R
	Switch to Monitoring view	Ctrl+Alt+M
	Switch to NAS view	Ctrl+Alt+N
	Switch to VMware view	Ctrl+Alt+N
	Toggle the Tools panel	Ctrl+Alt+T
	Cycle backward through panes	Shift+F6
	Cycle forward through panes	F6
	Previous screen	Alt+Left arrow
	Next screen	Alt+Right
	Save all changes	Ctrl+S
	Discard all changes	Ctrl+Z
	Refresh data	Ctrl+R
	Move to next item	Tab
	Move to previous item	Shift+Tab
	Open the Help Context menu	F1
Table Navigation	Move to the next row	Down arrow
	Move to the previous row	Up arrow
	Move to the next cell	Tab
	Move to the previous cell	Shift+Tab
	Leave table and move to the next item in the pane	Ctrl+Tab
	Leave table and move to the previous item in the pane	Shift+Ctrl+Tab

Location	Action	Shortcut
	Show context (right-click) menu for current table row	Shift+F10
Tree Navigation	Move to previous tree node	Up arrow
	Move to next tree node	Down arrow
	Collapse current tree node or move to parent of a collapsed node	Left arrow
	Expand current tree node or move to first child of an expanded node	Right arrow
	Show context (right-click) menu for selected tree node	Shift+F10
Tabs	Previous tab	Ctrl+Page Up
	Next tab	Ctrl+Page Down
Alarms	Show/hide Alarms panel	Ctrl+Alt+A
	Acknowledge All button	Ctrl+Shift+K

GUI Icons

Table 2. GUI Icons identifies the icons at the top of the GUI window.

Table 2. GUI Icons

Icon	Description	Shortcut
🔍	Find Objects — Checks for all matches against the names or descriptions for volumes, collections, pools, partners, and members, and displays the results in the Find Objects panel.	Ctrl+Shift+F
💾	Save Changes — Saves and applies any changes you made in a GUI window. If you do not save the changes, you are prompted to do so when you close the window or click another object in the tree.	Ctrl+S
↶	Discard Changes — Discards changes you made in a window.	Ctrl+Z
↻	Refresh — Refreshes the data that appears in the GUI. Do not use the browser refresh button to refresh the data that appears in the GUI.	Ctrl+R
◀ ▶	Navigate the GUI — Moves backward or forward through the GUI windows, according to the window history.	Alt+Left Arrow (to go back) Alt+Right Arrow (to go forward)
?	Help Menu — Opens the online help menu.	F1
▤ ▾	Tree view options — Opens a drop-down menu with various options for the different views (for example: Group, Volumes)	None

Search for Volumes, Collections, Pools, Partners, and Members

The search function checks for all matches against the names or descriptions for volumes, volume collections, storage pools, replication partners, and members in the group, and then displays the results in the Find Objects dialog box.

1. Click the magnifying glass icon or press Ctrl+Shift+F.

   The Find Objects dialog box opens.
2. Type the text that you want to search for.

About Online Help for Group Manager

In addition to tooltips and command-line help for the Group Manager GUI and CLI, online help is available for the Group Manager GUI. An Internet connection is required to use online help, which is served from a website in the Dell.com domain. You also have the option to install the help on your local system or a private web server.

Install the Group Manager Online Help as a Local Resource

When you launch online help for Group Manager, the default location is the following web server: psonlinehelp.dell.com/V.*release_number*

If you cannot access the Internet for any reason, you can download the online help kit and install it locally. Go to eqlsupport.dell.com to download the kit and installation instructions.

Access Online Help

Online help is available from the Group Manager GUI as either a set of help topics that you can search through, as a specific help topic, or as context-sensitive help that is built into a specific part of the GUI.

When you access the set of online help topics or a specific topic, the help opens in a separate browser window. To display the topics, either:

- Click the question mark icon at the top right of the Group Manager window and then select one of the following menu options:
  - **Help Contents** – Displays the table of contents for the set of help topics
  - **About *topic*** – Displays the help for that specific topic (for example, if you are in the Group window, the menu option is About Group and the topic displayed is About Groups)
  - **Keyboard Help** – Displays the shortcuts available from the keyboard
  - **About Group Manager** – Displays brief information about Group Manager
- Click the question mark icon in the lower right of the navigation pane.
- Expand the Tools tab at the bottom of the navigation pane and then click **Online help**.

From the browser window, you can use the table of contents or search option to find additional information

To access context-sensitive help about a specific GUI element, either:

- Click the embedded question mark icon.
- Click the **Help** button.

The system displays the help topic for the element in a separate browser window.

Localized Online Help

The Group Manager online help is available in the following languages:

- English (en)
- French (fr)
- German (de)
- Japanese (ja)
- Korean (ko)
- Simplified Chinese (zh)

- Spanish (es)

You can download non-English versions of the online help from eqlsupport.dell.com. The **readme** file in each language kit includes instructions for installing and using the localized online help.

Troubleshooting Online Help

If you have trouble launching the online help, check the browser security settings and the installed Java version.

Browser Security

Depending on your browser choice and the local Internet security settings, you might need to configure browser access to the help folder. For example, for Internet Explorer, you might need to add the help URL to the list of trusted sites:

1. Launch Internet Explorer, and go to **Tools**, then **Internet Options**.
2. Click the **Security** tab, then click **Trusted Sites** and then **Sites**.
3. Add the group's IP or management address to the list of trusted sites, using the format: `http://`*`group_ip_address`*.
4. Clear the option to require server verification, then click **Close**.
5. Click **OK** to close the Internet Options dialog box.

Java Version

Depending on the version of Java installed on your system, Internet Explorer might not display the online help locally if you run the Group Manager GUI as a browser applet. This difference is due to a regression in some versions of Java; you might not see the problem using browsers other than Internet Explorer.

Refer to the *Dell EqualLogic PS Series Storage Arrays Release Notes* for information about the Java versions supported for this release.

You must install and use the standalone Group Manager application as follows:

1. Do not close the browser window.
2. Click **Logout** to log out of the PS Series group.
3. Wait for the Group Manager welcome screen (`http://`*`group_ip_address`*`/welcome.html`). Do not select the option to require server verification.
4. Click **Launch as an application**. A Java window opens briefly, followed by the Downloading Application dialog box, which closes automatically on completion.
5. You can then log in to the group and set your user preferences to point to the local online help.

A Group Manager launch icon appears on your desktop, named for your PS Series group.

Internet Explorer Version

If you are using Internet Explorer version 10, you will need to switch off compatibility view in order to view help.

1. Launch Internet Explorer.
2. Click and select **Compatibility View settings**.
3. Clear the option to display websites in Compatibility view.

2

Architecture Fundamentals

The Dell EqualLogic product family provides a unified file and block storage platform.

Block-level storage consists of a sequence of bytes and bits of a certain length, called a block. Each block stores the data (like a hard drive) and the disk controller reads and writes data to the disks inside the storage array. Block-level access enables storage administrators to stipulate which block to send reads and writes to for the best performance. In file-level storage, the operating system keeps track of data in a directory of file and folder names. Access to data by an application is by file name and location within the file, and then translated into block-level access for physical reading and writing of data. Block-level transfers in a SAN are typically faster than file-level transfers in a NAS because they do not have any file system or network overhead.

The PS Series SAN provides block-level access to storage data through direct iSCSI access to PS Series arrays. The FS Series NAS appliance provides file-level access to storage data using NFS (UNIX) or SMB (Windows) protocols and the Dell FluidFS file system. This architecture uses a scale-out design, which enables you to increase performance as you add capacity, because the software manages the workload of data writes and the subsequent data reads across all the resources in the storage infrastructure.

PS Series Architecture

The PS Series architecture comprises three fundamental components: the PS Series array, the PS Series group, and the PS Series storage pool. These components coordinate operations and work together to form a virtual storage system.

The foundation of the PS Series architecture is a *PS Series array*. Each array is composed of redundant components: disks, controllers with mirrored write-back caches, network interfaces, power supplies, and cooling fans. When you configure a PS Series array, you add it to a PS Series *group*. A group can consist of up to 16 arrays of any family or model, as long as all arrays in the group are running firmware with the same major and minor release number. The group appears to the client servers as a single entity that offers network storage access in block mode.

NOTE: Dell recommends that all arrays run the same version of PS Series firmware at all times, except during the firmware update process.

Each array in a group is called a group *member*. You add members to the group, and allocate storage space in the group by creating storage *pools*. A pool can have from one to eight members, and you can create up to four pools in a group. Each member contributes its resources to the pool (disk space, processing power, and network bandwidth). As you add or remove members, the pool of storage space grows or shrinks. Hardware and other details remain hidden.

By default, a group member's RAID-protected disk space is added to the default storage pool. When you add a new member, it is automatically assigned to the default pool. The default pool cannot be deleted; however, you can create additional pools as required for your environment. You can assign members to the pools at any time.

The PS Series array software load balances data and network I/O to the group across all the group members' resources. When a group contains more than one member, those members might have different capacity levels and use different RAID levels. The load balancers transparently maintain optimal system performance and eliminate downtime to servers, applications, and users.

Next, you create volumes and assign them to the appropriate pool. (You can also organize volumes into folders for easy reference.) A *volume* provides the structure for the group. You can create a volume on a single group member or one that spans multiple group members. You identify a volume by specifying a unique volume name. Hosts on the network see these volumes as iSCSI targets. You set access controls for each volume so that only computers with an iSCSI initiator and the correct access credentials can access volume data.

Figure 2. PS Series Group and Pools depicts a PS Series group with three members and two storage pools. Table 3. PS Series Group and Pools explains the callouts used in the figure.

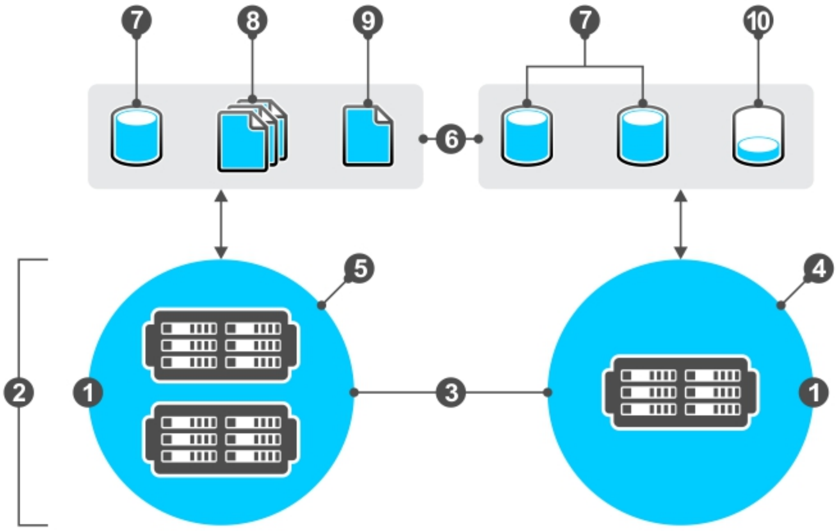

Figure 2. PS Series Group and Pools

Table 3. PS Series Group and Pools

Callout	Description
1	PS Series group Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block-storage devices.
2	PS Series members Each PS Series array is a member in the group and is assigned to a storage pool.
3	PS Series storage pools Containers for storage resources (disk space, processing power, and network bandwidth).
4	PS Series single-member storage pool A PS Series array represented as a member within a pool to which it is assigned.
5	PS Series multimember storage pool Multiple PS Series arrays represented as individual members within a pool to which they are assigned.
6	Storage space Space received from PS Series arrays to allocate data as needed through various structures (volumes, snapshots, thin provisioning, replicas, containers, SMB/NFS, quotas, and local users/groups).
7	Volumes Storage allocated by a PS Series group as addressable iSCSI targets.
8	Collection A set of volumes.

Callout	Description
9	Snapshots A point-in-time copy of data on a volume. Snapshots can be taken on a single volume or on a collection.
10	Thin-provisioned volume With thin provisioning, a minimal amount of space (10 percent by default) is reserved on a volume and then allocated when the space is needed.

FS Series Architecture

You can design a unified (block and file) storage architecture by adding a Dell FluidFS NAS appliance to a PS Series SAN. The FluidFS logical architecture integrates with the underlying SAN architecture. FluidFS presents a traditional file system to network clients while performing storage operations at the back end. This design utilizes all available resources at the network, server, and disk levels.

The FS Series architecture comprises four fundamental components: FluidFS NAS appliance, NAS cluster, NAS reserve, and NAS container.

The foundation of the FluidFS architecture is the FS Series *NAS appliance*. FluidFS software presents multiple network ports from multiple controllers to the client network. The system recognizes each individual NAS controller as a NAS member. Each NAS *member* can access and serve all data stored in the FluidFS system.

The *NAS cluster* is a virtual file server that hosts multiple SMB shares or NFS exports. You can have only one NAS cluster per PS Series group. NAS clients connect to file storage through one or more NAS client virtual IP addresses. The group IP address is never used by NAS clients to access the NAS cluster. The PS Series group stores and protects the block-based data. The NAS cluster serves the file-based data. When you configure a NAS cluster, you specify the network configuration for the NAS cluster and the amount of storage pool space that you want to allocate (the *NAS reserve* or available NAS datastore). You can increase, but not decrease, the NAS reserve. Dell FluidFS software, which runs on the NAS cluster, load balances client connections across the available NAS controllers in the NAS cluster.

To provision NAS storage, you need to create *NAS containers* within the NAS cluster. You allocate space for the NAS shares by creating these containers from the available space in the PS Series group pool. After you create the NAS containers, you can then configure a file-sharing protocol (NFS or SMB) to make storage data accessible over the network. NAS appears to a client as a file server and the client can map or mount drive shares. These shares and exports make storage space available to users.

Figure 3. PS Series Group with NAS Cluster depicts a unified storage architecture with both block and file storage. Table 4. PS Series Group with NAS Cluster explains the callouts used in the figure.

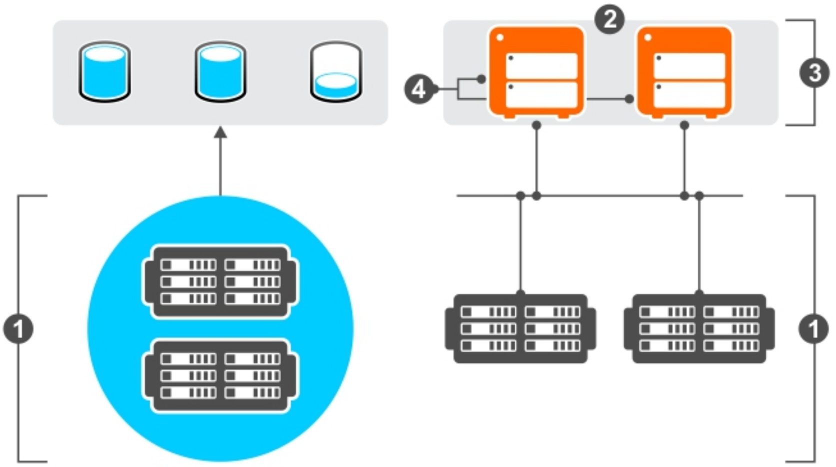

Figure 3. PS Series Group with NAS Cluster

Table 4. PS Series Group with NAS Cluster

Callout	Description
1	PS Series group Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block-storage devices.
2	NAS cluster Collection of NAS hardware (appliances) configured as part of a PS Series group. The FluidFS software runs on the cluster.
3	NAS appliances Hardware enclosures that contain NAS controllers.
4	NAS controllers Redundant, hot-swappable controllers in NAS appliances. The controllers interface over a fabric to the PS Series SAN storage.

About Groups

A PS Series group is a fully functional iSCSI storage area network (SAN).

You create a group when you configure one or more PS Series arrays and connect them to an IP network. In this virtualized storage system, the arrays become group *members* and share configuration data. A member belongs to a storage pool, and is configured with a specific RAID policy. Each member cooperates with other members to enable the virtualization of disk storage, controllers, caches, and network connections. Because the virtualization technology masks the underlying complexity of the storage configuration, client servers see the group as a single entity, giving you a centralized view of the storage data.

Figure 4. PS Series Group depicts PS Series groups. Table 5. PS Series Group explains the callouts used in the figure.

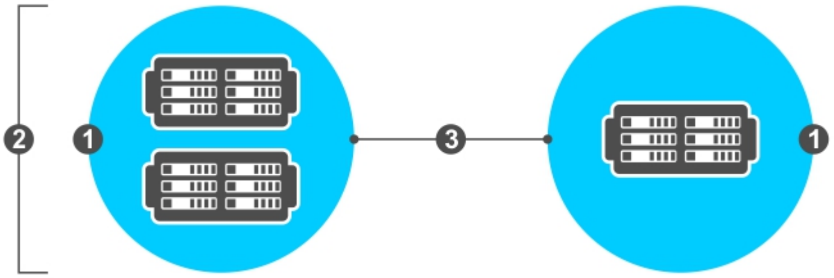

Figure 4. PS Series Group

Table 5. PS Series Group

Callout	Description
1	PS Series group Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block storage devices.
2	PS Series members One or more PS Series arrays represented as individual members within a pool to which it provides storage space to utilize.
3	PS Series storage pools Containers for storage resources (disk space, processing power, and network bandwidth). A pool can have one or more members assigned to it.

A group can provide both block and file access to storage data. Access to block-level storage requires direct iSCSI access to PS Series arrays (iSCSI initiator). Access to file storage requires the FS Series NAS appliance using NFS or SMB protocols and the Dell FluidFS scale-out file system.

With storage data management features, you can:

- Manage a group through several built-in mechanisms such as ssh, serial line, telnet, and web-based user interfaces. You do not need an external management station or management software.
- Configure the system to alert you to management activity or problems through log files, SNMP traps, and email notification
- Add more arrays (up to 16) to a group to increase capacity and performance
- Secure data and management access with authorization and authentication mechanisms
- Protect storage data with replication and snapshots

How Groups Work

Each group member cooperates with other members to automate resource provisioning and performance optimization. The storage system partitions the RAID-protected disk space that each member contributes to the storage pool into fixed-sized chunks of data, called pages. Pages logically separate the volume presented to the hosts from the physical resources of the storage array. Each volume has a page map that allocates pages to the members. The system automatically load balances data by performing transactional operations that move pages across member's disks and across all members in a pool. When a group receives a client server request, it identifies the location of the data and transfers the request to the member or members that contain the data.

As capacity and performance requirements increase, you can expand a group. When you add an array to an existing group, more space is immediately available. New members learn configuration and performance information from the group. You can also retire older equipment from the group as needed. You can choose a member to remove, and it will automatically offload its data to other members in the pool. The system automatically rebalances the load as the group scales. These operations are transparent to the servers, applications, and users.

Group: Configuration Recommendations

Before you configure a group, review the following recommendations.

- Make sure all the network interfaces on the members are configured, functioning, and accessible. If you have any issues, contact Dell Technical Support.

  NOTE: Limit configuration changes when a group has members that are offline. If you have members offline or a partitioned network, do not make configuration changes before you correct the problem to bring the group back to normal operation.

- Make sure the group has enough storage capacity. Low storage pool capacity generates an alert in SAN Headquarters. If a pool has less than 5 percent free space (or less than 100GB per member, whichever is less), a PS Series group might not have sufficient free space to efficiently perform the virtualization functions required for automatic optimization of the SAN. In addition, when storage pool free space is low, write performance on thin-provisioned volumes is automatically reduced to slow the consumption of free space. If pool capacity is low, try one or more of the following remedies:
  - Move volumes from the low-space pool to a different pool.
  - Increase pool member capacity by fully populating the drive bays or upgrading to higher-capacity disks.
  - Increase pool capacity by adding a member.
  - Reduce the amount of in-use storage space by deleting unused volumes or by reducing the amount of snapshot reserve.
- RAID policy

  Consider changing the RAID policy for a member. Change the policy only if you are sure that your applications will perform better with a different RAID level. For example, RAID 10 performs better than other RAID levels when a disk drive fails and when a RAID set is degraded. In addition, RAID policies with spare disks provide additional protection against drive failures.

  NOTE: You can change a RAID policy for a member only if the new RAID policy does not reduce the amount of available disk space. See Supported RAID Policy Conversions for more information.

- Volume management
  - Applications — Assign application volumes to a pool that includes members with the RAID policy that is optimal for the application.
  - Automatic load balancing — User fewer pools and let the group perform automatic performance load balancing.
- Hardware and firmware
  - To configure a hybrid array (model XS or XVS) to use RAID6 accelerated, the array must contain seven or more valid SSD drives at the time that the array is first set up.
  - Make sure member control module caches are in write-back mode.
  - Make sure all the group members are running the latest PS Series firmware.
- iSCSI connections

  Large, complex environments can use many iSCSI connections. A storage pool in a PS Series group can support numerous simultaneous connections, as described in the release notes for the particular Dell EqualLogic firmware release. These connections can be used for fully provisioned volumes, thin-provisioned volumes, snapshots, and control volumes (examples of control volumes are protocol endpoints and VDS/VSS).

  Attempting to exceed the supported number of connections results in an error message. You can reduce the number of iSCSI connections to the volumes and snapshots in a storage pool in several ways:

  - Disconnect from unused volumes and snapshots.
  - Modify MPIO settings to reduce the number of connections per volume.
  - Move volumes to another storage pool.
  - Create a new storage pool and move volumes to the new storage pool.
- Multipath I/O

  MPIO provides additional performance capabilities and network path failover between servers and volumes. For certain operating systems, the connections can be automatically managed. If MPIO is not creating multiple connections, you should:

- – Verify that the storage pool does not have the maximum number of iSCSI connections for the release in use.
- – Verify the access control policies for the volume. Using the iSCSI initiator name instead of an IP address can make access controls easier to manage and more secure.
- – Ensure that Dell EqualLogic MPIO extensions are properly installed on the supported operating systems. See the Host Integration Tools documentation for details.
- – Ensure that MPIO is supported and properly configured, according to the documentation for the operating system.

You can also monitor multiple PS Series groups with SAN Headquarters and can launch the Group Manager GUI from there; however, you cannot directly manage the storage from SAN Headquarters.

About Storage Pools

A storage pool is a container for a group member's storage resources. Storage pools allow you to allocate storage space into partitions based on the different types of storage resources and different types of data stored in the system.

Each member is assigned to a storage pool. A pool acts like a SAN within a SAN, creating an isolated storage environment within the overall PS Series SAN. Within the pool, load balancing happens automatically using available storage resources from the members in the pool.

You can design a homogeneous or heterogeneous pool.

- In a homogeneous pool design, all members of the pool are the same array model and use the same RAID level. In this case, the PS Series software automatically monitors data usage patterns and optimizes performance by the way it allocates data across the members.
- In a heterogeneous pool design, you can add members with different RAID levels or different models or a mix of both. In this case, the PS Series software uses load-balancing metrics to assign volumes to a RAID level appropriate for usage levels.

Single-Pool Design

Not all environments need multiple pools; a basic single-pool design might meet your needs. By default, each group has at least one pool, called the *default* pool. When you add a new member to a group, the system automatically assigns it to the default pool. You can rename the default pool, but you cannot delete it.

Multiple-Pool Design

Based on your business needs, you might want to divide the overall storage space into separate pools. A multiple-pool design allows you to prioritize applications within a SAN by placing them on separate storage resources, each optimally configured. Using this "SAN within a SAN," you can separate workloads as needed (for example: by application, service level, disk type, cost, or by department within the organization). As your capacity needs change, you can move members or volumes from one pool to another while data remains online.

NOTE: You can add both FS Series NAS reserve and PS Series volumes to a storage pool. However, you might want to keep the NAS reserve and the block (volume) space in different pools so that you can monitor space usage more easily.

Storage Pool: Design Checklist

Before you can design storage pools for your environment, you must identify your storage requirements: capacity, performance, data type, and applications. You can then make informed decisions about the storage pool design that best meets your needs.

Identify array specifications

- Disk type — Serial Attached SCSI (SAS) or Serial Advanced Technology Attachment (SATA)
- Disk size
- Disk speed

  NOTE: Some drive models do not carry information about their spin rate. For these models, the Group Manager GUI and CLI will show a speed of Unknown (or 0 rpm).
- RAID level

Identify requirements for the various types of data

- Requires 24/7 uptime and access
- For archival only
- Unique to specific departments (or example, the finance department might need exclusive access to certain data)

Identify application requirements

- List all applications accessing the data.
- Calculate the disk space, network bandwidth needs, and performance characteristics for each application. For example, some applications require many random data transfers, while others require only a few large sequential data transfers.
- Identify applications that require priority access to data.

Other requirements

- Each member can be assigned to only one pool.
- Each pool can host up to eight members.
- Each group will have at least one pool (the default pool).
- Each group can have up to four pools.
- To use multiple pools, a group must have at least two members.
- You can add or move members to different pools.

Storage Pool: Design Examples

You can create a single-pool or multiple-pool design for your configuration. The size of the PS Series group — the number of members and their capacity — can help determine how you plan the storage.

Example 1: Use pools to isolate crucial applications or data.

- You can create a separate pool and set the access control so that only select servers have access to that pool. For example, if you have mission-critical data to protect, you might create a storage pool that contains the highest-performance arrays in your environment.
- You might want to separate application data from backup and archive data, or you might want more control over which volumes run on which hardware. In this case, you can assign each volume to a different storage pool where each pool is independent of the other.

Example 2: Use pools to segment data by organization or type of data.

- You can segment storage by business units. For example, you can create a separate pool for each business unit. By assigning each business unit's data to a separate pool and assigning access controls, you can ensure that business units have exclusive use of their own storage.
- You can segment bulk storage for long-term retention. You can create a separate pool with arrays configured to maximize storage capacity.

Example 3: Use pools to optimize performance of your data.

- You might want to optimize the performance of the database application. You can implement three pools: one for the database, one for applications, and one for backup and archival data. In this case, the database would be assigned to the pool that contains the members with the best random-access performance.

Tip: Dell EqualLogic recommends using as few pools as possible and letting the PS Series software handle load balancing and assignment of volumes to the most appropriate storage.

About Volumes

Volumes provide the storage allocation structure within the PS Series group.

To access storage in a PS Series group, you allocate portions of a storage pool to volumes. You can create a volume on a single group member or one that spans multiple group members. You assign each volume a name, size, and a storage pool. The group automatically load balances volume data across pool members.

Figure 5. PS Series Volumes depicts volumes in a PS Series group. Table 6. PS Series Volumes explains the callouts used in the figure.

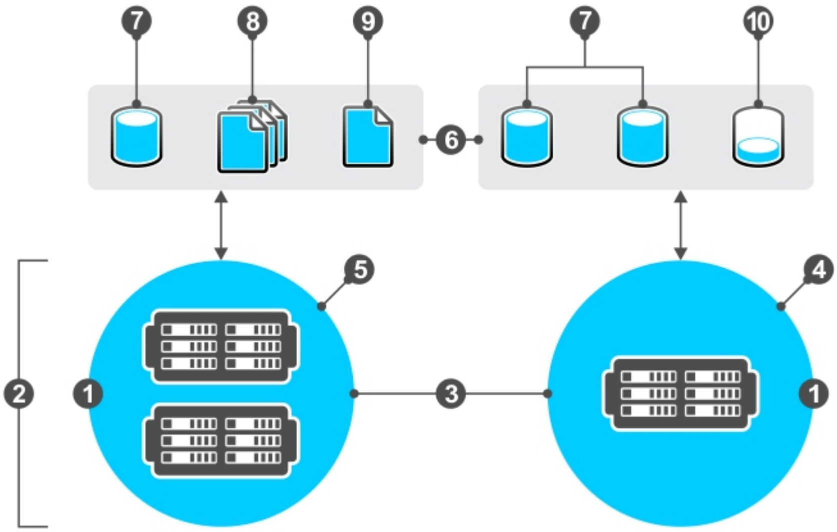

Figure 5. PS Series Volumes

Table 6. PS Series Volumes

Callout	Description
1	PS Series group Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block storage devices.
2	PS Series members Each PS Series array is a member in the group and is assigned to a storage pool.
3	PS Series storage pools Containers for storage resources (disk space, processing power, and network bandwidth).
4	PS Series single-member pool A PS Series array represented as a member within a pool to which it is assigned.
5	PS Series multimember pool Multiple PS Series arrays represented as individual members within a pool to which it is assigned.
6	Storage space

Callout	Description
	Space received from PS Series arrays to allocate data as needed through various structures (volumes, snapshots, thin provisioning, replicas, containers, SMB/NFS, quotas, and local users and groups).
7	Volumes Storage allocated by a PS Series group as addressable iSCSI targets.
8	Collection A set of volumes.
9	Snapshots A point-in-time copy of data on a volume. Snapshots can be taken on a single volume or on a collection.
10	Thin-provisioned volume With thin provisioning, a minimal amount of space (10 percent by default) is reserved on a volume and then allocated when the space is needed.

For each volume, the group generates an iSCSI target name, which you cannot modify. An iSCSI target name includes a prefix, a string, and the volume name. Initiators use the target name to connect to a volume. For example:

iqn.2001-05.com.equallogic:7-8b0900-6d0000000-001ebbc5d80sf0k0-db3

In this example:

prefix is	iqn.2001-05.com.equallogic
string is	7-8b0900-6d0000000-001ebbc5d80sf0k0
volume name is	db3

Each volume appears on the network as an iSCSI target. Hosts with iSCSI initiators use the volume's target name to connect to the volume.

Each iSCSI volume supports a set of features and capabilities:

- Snapshots – To protect volume data from mistakes, viruses, or database corruption, you can use snapshots.
- Replication – To protect against disasters, you can replicate volume data from one group to another.
- Thin Provisioning – To manage storage capacity utilization on demand, you can use thin provisioning.
- Clones – To create a master or boot image, full system backup, or transfer a system to another person, you can use cloning.
- Volume Unmap (block reclaim) – To recover space previously allocated to volumes, you can unmap them.
- Volume Undelete – To restore mistakenly deleted volumes, you might be able to use volume undelete.

  NOTE: The system permanently deletes volumes after 7 days, and sometimes sooner if the space is needed.
- Volume Folders – To organize volumes into folders for quick visual reference, you can use volume folders. Using folders has no effect on the contents of the volumes.
- Control Access to iSCSI Initiators – To protect your volumes from unauthorized and uncoordinated access by iSCSI initiators, you can use access control policies.
- Control Access to Hosts (servers) – To prevent inadvertent corruption of the volume caused by multiple hosts writing to it in an uncoordinated manner, enable multihost access to a volume.

Volume Attributes

You set some attributes when you create a volume; other attributes use default values. In most cases, you can modify all the volume attributes. Template volumes and thin clones have some restrictions.

Table 7. Volume Attributes describes the attributes that allocate space and set the characteristics of a volume.

Table 7. Volume Attributes

Volume Attribute	Description
Name	Volume name is unique in the group. The volume name appears at the end of the iSCSI target name, which the group generates automatically. Computer access to the volume is always through the iSCSI target name, rather than the volume name.
Description	Optional description for the volume — up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
Storage pool	Name of pool for the volume. The group stores all the volume data on the pool members. The default is the default pool. Thin clones must reside in the same pool as the template volume. If you move a template volume to a different pool, all the attached thin clones also move.
Reported size	Reported size of the volume in MB, GB, or TB. The group rounds up volume sizes to the next 15MB if the size is not a multiple of 15MB. You cannot change the reported size of a template volume.
Thin provisioning settings	Controls whether the volume is thin provisioned and, if so, the minimum and maximum volume reserve and the in-use space warning limit. The defaults are no thin provisioning and the groupwide volume settings.
Snapshot reserve	Optional amount of space to reserve for snapshots of the volume, based on a percentage of the volume reserve. The default is the groupwide snapshot reserve setting. If data reduction has been enabled on the volume, snapshot reserve is permanently disabled.
iSCSI alias	Name that some iSCSI initiators display. Use the alias to identify the iSCSI target. The default is the groupwide volume setting.
Permission	Indicates whether the volume is read-write (the default) or read-only. You cannot set a template volume to read-write permission.
Administrative status	Indicates whether the volume is online (the default) or offline. Initiators cannot discover or connect to an offline volume.
Access controls	Conditions that hosts must meet to access the volume and its snapshots. To allow volume or snapshot access, you must create at least one access control policy. You can create the policy when you create a volume or after you create the volume.
Administrator	You can assign a volume to a specific volume administrator.
Multihost access setting	Indicates whether the volume allows or disallows (default) access from initiators with different IQNs.
iSNS discovery setting	By default, iSNS servers cannot discover iSCSI targets in a group. To allow discovery by iSNS servers, you must enable this functionality on a volume or snapshot.
RAID preference	RAID policies are established on a per-member basis. You can override this policy by enabling a RAID level preference for individual volumes. Thin clones inherit the RAID preference of the template volume.
Sector size	Size of the sector. The sector size of a volume does not depend on the sector size of the physical drives in the array. The default volume setting of 512 bytes is almost always optimal, even if the physical drives use 4K-byte sectors. You cannot change the sector size after creating the volume.

Volume Types

A PS Series group supports the following volume types:

- Standard

  The default volume type is a standard volume. No restrictions apply to a standard volume. You can enable (and disable) thin provisioning on a standard volume.
- Template

  A template volume is a type of volume that is useful if your environment requires multiple volumes that share a large amount of common data. After you write the common data to a standard volume, you can convert it to a template volume and then create thin clones. Template volumes are read-only to protect the common data.
- Thin clone

  Thin clones are based on a template volume and enable you to use space efficiently in storage environments that require multiple volumes with a large amount of common data. After you create a thin clone, you can write to the thin clone as needed.

You can replicate any volume type, resulting in a replica set for the volume. In addition, you can fail over any volume type, resulting in a recovery version of the volume. However, you can only fail back a standard volume or a thin clone volume.

About Volume Space Allocation

To size volumes correctly, you need to understand how the group allocates space to volumes. Although you can modify a volume size, some operating systems and initiators do not easily handle size changes.

When you create a volume, you specify the reported size for the volume, which is the maximum amount of space that the group might be able to allocate to the volume. You can increase or decrease the reported size as needed.

The reported size is seen by iSCSI initiators. If a write to a volume exceeds the reported size, the write fails, and the group generates event messages.

The actual amount of pool space that the group allocates to a volume is called the volume reserve. The value of the volume reserve depends on whether you enable thin provisioning on a volume:

- Thin provisioning disabled

  The volume reserve is equal to the reported size.
- Thin provisioning enabled

  If you enable thin provisioning on a volume, the group allocates space based on volume usage. The volume reserve is equal to or less than the reported size, depending on volume usage and the thin-provisioning settings.

Space allocated for volume operations (for example, snapshot reserve and local replication reserve) is based on the volume reserve.

You cannot use space that the group allocates to a volume (or for volume operations) for other purposes. Therefore, make sure you allocate space only when necessary.

You must fully understand application and workload space requirements to allocate the correct amount of space.

About NAS Architecture

NAS architecture is an advanced clustered architecture, providing the performance of a networked file system with the efficiency of a storage system.

Figure 6. NAS Hardware Architecture shows a typical configuration, although other types of configurations are possible. Table 8. NAS Hardware Architecture explains the callouts used in the figure.

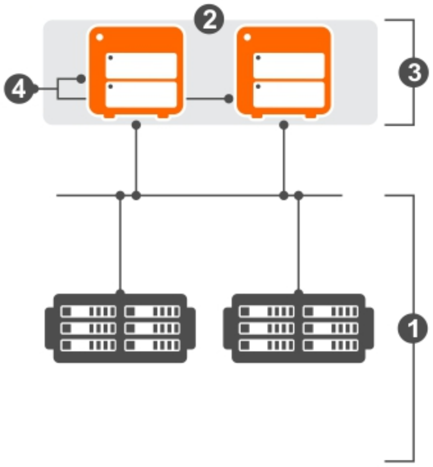

Figure 6. NAS Hardware Architecture

Table 8. NAS Hardware Architecture

Callout	Description
1	PS Series group (partial) Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block-storage devices.
2	NAS cluster Collection of NAS hardware (appliances) configured as part of a PS Series group.
3	NAS appliances Hardware enclosures that contain NAS controllers.
4	NAS controllers Redundant, hot-swappable controllers in NAS appliances. The controllers interface over a fabric to the PS Series SAN storage.

Figure 7. NAS Software Architecture shows a logical representation of the NAS cluster. Table 9. NAS Software Architecture explains the callouts used in the figure.

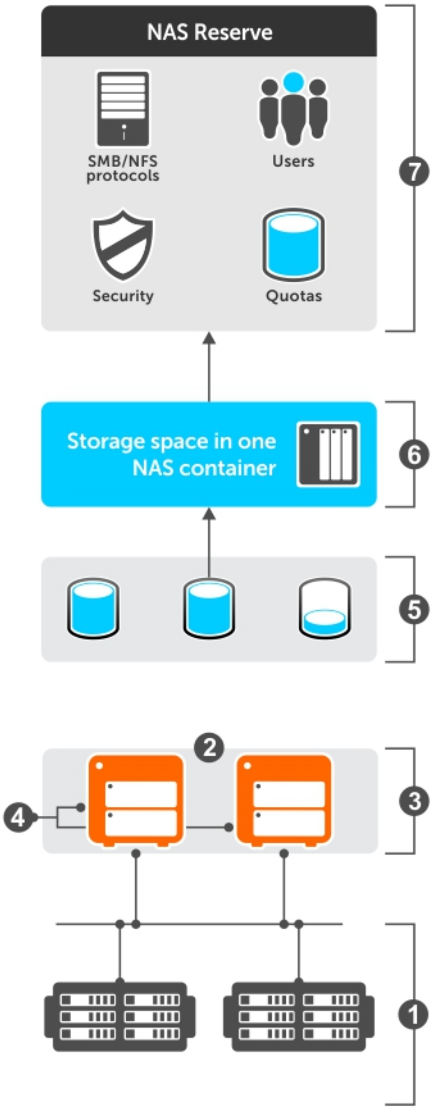

Figure 7. NAS Software Architecture

Table 9. NAS Software Architecture

Callout	Description
1	PS Series group (partial) Storage area network (SAN) comprising one or more PS Series arrays connected to an IP network. Arrays are high-performance (physical) block-storage devices.
2	NAS cluster Collection of NAS hardware (appliances) configured as part of a PS Series group. The FluidFS software runs on the cluster.
3	NAS appliances Hardware enclosures that contain NAS controllers.
4	NAS controllers

Callout	Description
	Redundant, hot-swappable controllers in NAS appliances. The controllers interface over a fabric to the PS Series SAN storage.
5	NAS containers Total amount of available storage space allocated for NAS shares
6	Storage space Space that allocates data as needed through various structures (volumes, snapshots, thin provisioning, replicas, containers, SMB/NFS, quotas, and local users and groups)
7	NAS reserve Amount of available storage space allocated to the NAS cluster for storing internal data and user data. This data includes: • SMB/NFS protocols – SMB shares provide users a way to share files and data across a Windows network, while NFS exports provide users a way of sharing files and data across UNIX networks. NFS clients can only mount exported directories. • Local users and groups – Individual accounts from which users can access SMB/NFS shares. These accounts can be grouped so that they share the same access permissions. • Quotas – Define how storage space on a NAS container is allocated among users and groups of users. • Security – Volume-level and group-level access controls for NAS containers.

About NAS Clusters

NAS clusters enable integration of file storage and block storage. A NAS cluster contains hardware with two redundant, hot-swappable controllers and RAID protection for internal-use disks. The scalable Dell Fluid File System (FluidFS) enables PS Series storage arrays to store NAS cluster and client data.

A NAS cluster is made up of NAS appliances, configured and networked as described in your NAS appliance hardware documentation.

NOTE: Only an account with group administrator privileges can view or manage a NAS cluster.

When you configure a NAS cluster, you specify the network configuration for the service and the amount of storage pool space for the NAS reserve. The NAS reserve is configured with Fluid FS and stores client data, in addition to NAS cluster metadata. The NAS reserve is divided into NAS containers in which you create SMB shares and NFS exports to make storage space available to users.

NOTE: Depending on the NAS appliance model, the NAS member might show its controllers organized under the NAS member as a NAS controller pair.

Both NAS controllers in a NAS controller pair operate simultaneously. Each NAS controller has a high-performance cache that is mirrored by its peer NAS controller. If one NAS controller in a NAS controller pair fails, the other NAS controller automatically continues operation with no impact on availability.

Just as you can add arrays to a PS Series group to expand SAN capacity, you can start with one NAS appliance and then add another NAS appliance to the NAS cluster as a NAS controller pair. Adding a second NAS controller pair increases NAS cluster performance.

You can create multiple NAS containers in a NAS cluster. NAS containers have robust security mechanisms and support snapshots and NDMP for data protection. On each NAS container, you can create multiple SMB shares and NFS exports. Clients with the correct credentials can access the shares and exports.

Clients connect to NAS containers through a single NAS cluster IP address, providing a single-system view of the NAS storage environment. For availability and performance, client connections are load balanced across the available NAS controllers.

A NAS cluster can serve data to multiple clients simultaneously, with no performance degradation. Clients connect to NAS storage through the NAS protocols of their operating system:

- UNIX users access NAS storage through the NFS protocol.
- Windows users access NAS storage through the SMB protocol.

After the client establishes the preliminary connection, the NAS storage acts as a normal storage subsystem, accessed in the usual way by users or applications.

The NAS cluster hardware and network configuration is described in detail in your NAS appliance hardware documentation and includes the following components:

- NAS controller pair
- Client network

  Used for client access to the NFS exports and SMB shares hosted by the NAS cluster. For security, the client network is typically separate from the SAN and internal network.
- SAN/internal network

  Enables internal communication between the controllers as required for failover, communication between controllers, and communication between the NAS cluster and the PS Series SAN. The SAN/internal network is connected to the same set of switches. For performance and security reasons, Dell strongly recommends that this network be in a different subnet from the client network.

A NAS cluster supports various network topologies. From a performance perspective, it is important to consider the subnets to which the clients belong (they might belong to more than one) and the subnets to which the NAS cluster belongs.

For example, if the NAS cluster and all clients are on the same subnet, the network is considered “flat,” which provides the best performance.

If clients reside on a subnet that is not the same as the NAS cluster client subnet, the network is routed, and clients can access NAS container data by using a router or layer 3 switches. In a routed network configuration, you should configure multiple NAS cluster IP addresses for proper load balancing.

Network interfaces in a NAS controller are numbered. Each interface has a specific function and is used for a client network, SAN network, internal network, or IPMI connection. If you understand the function of each network interface, you can ensure a highly available network configuration in which no single switch failure or network cable disconnection results in a service disruption.

About NAS Controllers

A NAS controller has preinstalled NAS firmware and redundant, battery-backed, hot-swappable hardware. NAS controllers do not require day-to-day maintenance. The group logs alarms and events related to the NAS controllers.

When a NAS appliance is functioning properly, both NAS controllers serve NAS container data. Data caching across controllers provides high performance, and cache mirroring between controllers ensures data redundancy.

Each NAS appliance contains two controllers that handle client connections, manage read and write disk operations, and interact with servers and workstations. I/O load-balancing mechanisms direct client requests to the least busy NAS controller, which maintains an even load balance across all NAS controllers in the NAS cluster.

If one controller in a NAS appliance fails, clients fail over automatically to the peer NAS controller, and cache mirroring stops (journaling mode). When failover occurs, some SMB clients automatically reconnect to the peer NAS controller. In other cases, an SMB application might fail, and you must restart it. NFS clients experience a temporary pause during failover, but client network traffic resumes automatically.

NOTE: Not all NAS cluster management operations are supported while a NAS controller is down or detached. Therefore, it is important to replace a failed NAS controller as soon as possible.

You can perform the following NAS controller operations:

- Add a NAS appliance to a NAS cluster to increase processing power and allow more client connections
- Replace a failed controller

  If a NAS controller fails, the NAS cluster is still operational, but you cannot perform most service configuration modifications until you detach the failed NAS controller. While a NAS controller is down or detached, performance might decrease because data is no longer cached.
- Rebalance or fail back client connections after changing the hardware configuration (for example, adding a NAS appliance or attaching a NAS controller)

  You can rebalance client connections across all the functioning NAS controllers, or you can fail back client connections to recently recovered NAS controllers.
- Update the firmware on a NAS controller
- Cleanly shut down a NAS controller (for example, if you are moving the NAS cluster hardware)

You can perform some NAS controller maintenance operations without impacting NAS member operation (for example, replacing a power supply). Other NAS member maintenance operations require shutting down the NAS controller before performing maintenance procedures and then restarting the NAS controller.

About NAS Containers

To provision NAS storage, you can create multiple NAS containers in a NAS cluster. In a NAS container, you can create multiple SMB shares and SMB home shares and NFS exports. Access to shares and exports is through one or more NAS cluster virtual IP addresses.

The number and size of the NAS containers in a NAS cluster depend on the storage needs of your NAS clients and applications. You can increase or decrease the size of a NAS container as needed without disruption to the NAS clients accessing the NAS container.

In addition, NAS containers:

- Have robust security mechanisms
- Support user and group quotas
- Support snapshots for data protection
- Support thin clones
- Support NDMP for remote backups
- Support replication to remote FS Series NAS clusters for disaster tolerance
- Support thin provisioning

You can create a single, large NAS container, or you can create many NAS containers. Creating multiple NAS containers enables you to apply different management policies to the NAS containers, as required by your organization. For example, you can apply different backup, snapshot, security, and quota policies to the NAS containers.

When you create a NAS container, SMB share, SMB home share, or NFS export, the NAS cluster applies default values.

About the NAS Reserve

When you configure a NAS cluster, you allocate a portion of storage space to the NAS reserve. As part of the service configuration, the NAS reserve is configured with the Dell Fluid File System (FluidFS) and Dell EqualLogic Auto-Snapshot Manager/Microsoft Edition (ASM/ME).

Figure 8. NAS Reserve depicts the components of the NAS reserve. Table 10. NAS Reserve explains the callouts used in the figure.

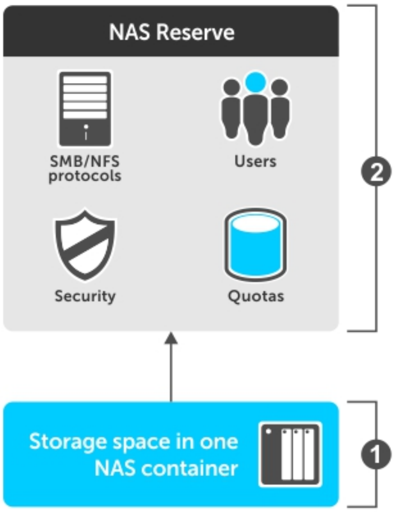

Figure 8. NAS Reserve

Table 10. NAS Reserve

Callout	Description
1	NAS storage space Space allocated for storing user data as needed through various structures (volumes, snapshots, thin provisioning, replicas, containers, SMB/NFS, quotas, and local users and groups)
2	NAS reserve Amount of available storage space allocated to the NAS cluster for storing internal data and user data. This data includes: • SMB/NFS protocols – SMB shares provide users a way to share files and data across a Windows network, while NFS exports provide users a way of sharing files and data across UNIX networks. NFS clients can only mount exported directories. • Local users and groups – Individual accounts from which users can access SMB/NFS shares. These accounts can be grouped so that they share the same access permissions. • Quotas – Define how storage space on a NAS container is allocated among users and groups of users. • Security – Volume-level and group-level access controls for NAS containers.

You can specify the total size of the NAS reserve and NAS container space, and the system will automatically add the required internal capacity. The system deducts a certain amount of the NAS reserve for internal use. The exact amount of internal space varies by configuration (see Table 11. Minimum and Maximum NAS Reserve Capacities), but it is roughly calculated as a fixed amount of space for each controller pair, plus approximately 0.5 percent of the total NAS reserve.

The minimum and maximum capacities for the NAS reserve vary based on the cluster's configuration. You can increase the size of the NAS reserve as your NAS storage space requirements increase. However, you cannot decrease the size of the NAS reserve.

Table 11. Minimum and Maximum NAS Reserve Capacities

Cluster Configuration	Minimum NAS Reserve	Maximum NAS Reserve
2 controllers	512GB	510TB
4 controllers	1024GB	510TB

3

Set Up the iSCSI SAN

To start using the PS Series array:

1. Configure the array on the network and create a PS Series group. See the *Installation and Setup Guide* for more information.
2. Log In to the Group Manager GUI.
3. Set the RAID Policy and Pool for a New Member (and assign the member to the default pool).
4. Create a Volume.
5. Connect Initiators to iSCSI Targets.

4

Post-Setup Tasks

After you complete the initial setup and deployment of your PS Series array, Dell strongly recommends that you perform certain tasks to finish configuring the group.

Perform the following important tasks to fully configure the group:

- Set the group time and date
- Customize access to the group GUI and CLI
- Configure session settings
- Set up event notifications
- Set groupwide volume defaults
- Secure your data
- Add the group to SAN HQ

When you need to update the firmware version that is running on an array, see *Updating Firmware for Dell EqualLogic PS Series Storage Arrays and FS Series Appliances*. This document is available from the support site, eqlsupport.dell.com.

About the Group Date and Time

All members of a group share the same time zone. Each array's clock is set at the factory, based on GMT. The default time zone is America/New York, ET (Eastern Time).

Group time is initially based on the time set on the array that was used to create the group. Each array you add updates its clock to match the group's date and time. The group date and time determines the timestamp used to identify some objects (snapshots, for example) that you create on the group. Changing the time is a groupwide operation and updates all members.

You can change the group time manually, or you can configure an external Network Time Protocol (NTP) server to automatically set the same time for all members.

Change or Delete an NTP Server

A PS Series group can use up to three external Network Time Protocol (NTP) servers to automatically set the same time for all the group members.

The following considerations apply:

- If you are using a dedicated management network, Dell recommends that your NTP server be on the same subnet as the dedicated management network.
- If the group time changes by an hour or more, either because you manually modified it or the NTP server time changed, the synchronization operation can take as long as 24 hours.
- If you have a NAS cluster in your PS Series group and are using Active Directory for external authentication, the NTP server must be the same NTP server used by Active Directory. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for more information about NTP configurations.

To change or delete NTP servers:

1. Click **Group** → **Group Configuration**.

2. Click the **General** tab.
3. Click **Add** under NTP servers in the Date and Time panel.
4. Type the IP address for the NTP server.
5. Type the port number for the NTP server.
6. Click **OK**.
7. Select the IP address and click **Modify** or **Delete** as needed.
8. Use the arrows to move a server up or down in the list.

Change the Time Zone and Clock Time

To display the current date and time values:

1. Click **Group** → **Group Configuration**.
2. Click the **General** tab.
3. Select the time zone and city from the **Time zone** drop-down list.
4. To set the time, click **Change** to open the **Change group time** dialog box. You can either:
   - Type the date and time in the **New date and time** field
   - Synchronize the time with the computer running the GUI by clicking **Set to time of system running GUI**
5. Click **OK** to save the changes. Click **Cancel** or **Reset to current group time** to cancel the changes.

About Session Idle Timeout

Group Manager can automatically log itself out after a specified idle period. When this option is enabled, you can specify a session idle timeout in minutes, with a minimum value of 1 minute. If a Group Manager session detects no user activity after the session idle time has elapsed, it displays a countdown progress bar. When the countdown progress bar appears, you have 1 minute to resume or log out of the session before you are automatically logged out.

NOTE: Setting or changing the logout time in the Group Manager GUI takes precedence over changes made using the CLI.

Group Manager can log out of the following session types:

- Group Manager GUI
- Group Manager CLI using SSH, telnet, or console
- FTP

Change the Session Idle Timeout

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. Select the **Session idle timeout** checkbox, which enables the session idle timeout function.
4. Select a time (in minutes) from the counter. The minimum time that you can specify is 1 minute. The maximum time that you can specify is 1440 minutes (24 hours). The default value is 30 minutes.
5. Click **Save** at the top of the Group Manager window.

To disable session idle timeout, clear the **Session idle timeout** checkbox. (You cannot set the timeout value to zero.)

Set General GUI Policies

You can set local GUI policies such as the display font size.

1. Click **Tools**.
2. Click **User Preferences**. The **Modify user preferences** dialog box opens, displaying the **General** tab.

3. Select your preferences and click **OK**.

Set GUI Communication Policies

You can set policies for managing connections between your workstation and the Group Manager GUI.

1. Click **Tools**.
2. Click **User Preferences**.
3. Click the **Communication** tab.
4. Select your preferences and click **OK**.

Set Alarm Policies

You can set alarm policies to control problem notification.

1. Click **Tools**.
2. Click **User Preferences**.
3. Click the **Alarms** tab.
4. Select your preferences and click **OK**.

An alarm indicates a potential problem with your array.

Display NAS Cluster Alarms

NAS cluster alarms help you discover and correct problems before they disrupt operations. You can display alarms by severity or by category.

NOTE: A problem with a NAS controller does not always generate an alarm, but it will generate an event message.

Display by Severity

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **View alarms**.

Display by Category

1. Click the **Up arrow** in the corner of the Alarms panel, which is located at the bottom of the Group Manager GUI window, to open the Alarms panel.
2. Click the **Critical** tab to display the critical alarms. Click the **Warnings** tab to display the warning alarms.

Set Advanced Policies

You can select options for setting data validation, debugging policies, and software updates. You can also set advanced policies to control problem notifications.

1. Click **Tools**.
2. Click **User Preferences**.
3. Click the **Advanced** tab.
4. Select your preferences and click **OK**.

Monitor Events

Event messages help you monitor group operations and correct problems before they disrupt operations. The group generates a message when a significant event that requires corrective action occurs in the group (for example, when hardware fails or replication space is insufficient). The group also generates event messages when certain normal operations occur (for example, when a user logs in to the group or creates a volume).

To display events:

1. Click **Monitoring**.
2. Under Events, select **Event Log**.

   The events display in the window. To change which group's events display in the window, select the group from the **View** drop-down menu.

From the Event Log window, you can:

- Display all events or events of a specific priority. Use the **View** menu to select the events that you want to display (All events, Warnings and errors, Errors, and NAS events).
- Retrieve previous events. To retrieve the most recent 100 events, click the **More** button. Click it again to retrieve the next 100 events.
- Acknowledge all events. Unacknowledged events appear in bold. To acknowledge the receipt of all event messages, click the **Acknowledge all** icon (✓).
- Clear the event list. To erase all the events from the panel, click the **Clear event list** icon (✖). To show the events again, click **More**.
- Show or hide event details:
  - Move the pointer over an event. A pop-up window opens, showing event details.
  - Double-click an event. The event details panel opens at the bottom of the events list.
  - Select an event and click the **Show/Hide event details** button near the upper-right corner of the window. The event details panel opens or closes at the bottom of the events list.

  If you are showing details, the following information is displayed:

  - Level — Event type
  - Time — Date and time the administrator action took place
  - Member — Group member on which the action was performed
  - Subsystem — Array subsystem in which the action took place
  - Event ID — System identifier for the event
  - Message — Description of the event

Monitor Audit Log Events

The audit log shows events that are related to the accounts on the group. The events are listed in reverse chronological order, with the most recent entries displayed first.

To display audit log events:

1. Click **Monitoring**.
2. Under Events, select **Audit Log**.

   The audit log events display in the window. To change which group's events display in the window, select the group from the **View** drop-down menu.

From the Audit Log window, you can:

- Display all events or events of a specific group account. Use the **View** menu to select the account for which you want to display events:
- Retrieve previous events. To retrieve the most recent 100 events, click the **More** button. Click it again to retrieve the next 100 events.
- Acknowledge all events. Unacknowledged events appear in bold. To acknowledge the receipt of all event messages, click the **Acknowledge all** icon (✓).

- Clear the event list. To erase all the events from the panel, click the **Clear event list** icon (❌). To show the events again, click **More**.
- Show or hide details about a specific event:
  - Move the pointer over an event. A pop-up window opens, showing event details.
  - Double-click an event. The event details panel opens at the bottom of the events list.
  - Select an event and click the **Show/Hide details** icon near the upper-right corner of the window. The event details panel opens or closes at the bottom of the events list.

  If you are showing details, the following information is displayed:

  - Level — Event type; for audit log events, the type is always Audit
  - Time — Date and time the administrator action took place
  - Member — Group member on which the action was performed
  - Subsystem — Array subsystem in which the action took place
  - Event ID — System identifier for the event
  - Message — User action or audit event

About Event Notifications

Dell recommends that you configure event notification so that you automatically receive messages when certain events occur in the group (and NAS cluster, if configured). Events enable you to track operations and also detect and solve problems before they affect performance or data availability. Notifications are not enabled by default, and must be configured manually.

Select one or more notification methods to automatically receive notifications when events occur:

- Email notification — If an event occurs, the group automatically sends a message to designated email addresses.
- Email Home — If a hardware component fails or if you update firmware, the group automatically notifies customer support. Email Home is available to all customers, but response time and assistance is based on the validity and level of your support contract.
- Remote server logging — The group logs events to a remote server at the syslog facility. You can also access events from the syslog server.
- Software Update notifications —The group periodically checks the customer support website for updates to Dell EqualLogic software, including PS Series array firmware.

About Email Notification

If an event occurs, the group automatically sends a message to designated email addresses.

The group collects multiple events into a single message, eliminating the need for multiple emails. If only one event occurs within 1 minute, the group sends email to the addresses that you configured for notification. If another event occurs within 1 minute, the timer starts over and sends email after 2 minutes.

The following prerequisites apply:

- To use email notification, a group must have access to a Simple Mail Transfer Protocol (SMTP) server or email relay.
- Up to five email addresses to receive notifications.
- An IP address and optional port for the SMTP server in the format `ip_address:port`. The default port is 25.

  You can enter up to three IP addresses. The group uses one SMTP server or email relay at any time. The first server that you specify is the default server. The group uses the other servers in the order specified, if the default server is not available.
- An email address that will appear in the *From* field in the notification email. You can use your own email address or the group name at your company's email address (such as `GroupA@company.com`).

  When the intended recipient receives email, the email itself specifies which group it came from. This information is helpful in multigroup environments, and reduces the chance that the email server or recipient discards or rejects notifications.

Configure Email Notifications

You can define the list of email recipients to whom notifications will be sent for various alert levels. You can also change the email notification configuration at any time.

1. Click **Group** → **Group Configuration**.
2. Click the **Notifications** tab to open the Email Event Notifications panel.
3. If it is not already selected, select the **Send email to addresses** checkbox.
4. In the Email recipients section, click **Add** to open the dialog box. Type an email address and then click **OK**.

   You can specify up to five email addresses.
5. To choose the types of events that will generate email notifications, select or clear the checkboxes in the **Event priorities** section.
6. In the **SMTP servers** section:

   a. Click **Add** to open the Add SMTP server dialog box.

      You can configure up to three servers, and you must have at least one server configured to send email.

   b. In the **IP address** field, enter the IP address for the SMTP server.

      The format is *A.B.C.D*, where A, B, C, and D are 1- to 3-digit numbers.

   c. In the **Port** field, type the port number to use to communicate with the specified server, or click the **Use default port** button.

      The default port, 25, is used if no port is defined.

   d. Repeat steps a through c to add more servers.

   e. To change the order of the listed servers, select an IP address and click the **Up** or **Down** link to reposition the address in the list.

      Servers are listed in the order that they should be used for email.
7. Click **OK** and then type an email address in the **Sender email address** field.

Test the Email Notification Configuration

1. In the Event Logs panel, select **Enable automatic delivery of INFO messages**.
2. Log out of the group and then log back in to the group.

If an email recipient does not receive notification of the logout and login events, verify and correct the configuration.

About Email Home

NOTE: Support for Email Home will be discontinued in a future release. Email home will be replaced by SupportAssist, which provides functionality similar to Email Home. Dell recommends installing SAN Headquarters version 3.2 (or later) and enabling SupportAssist to monitor your PS Series groups and report and upload critical events to Dell Technical Support. You do not have to wait until the next release to switch from Email Home to SupportAssist.

In the current release, if a hardware component fails or if you update firmware, the group can automatically notify customer support through email. Dell recommends that you enable Email Home to expedite Dell Technical Support becoming engaged in solving any problems.

Email Home is available to all PS Series customers, but response time and assistance is based on the terms of your support contract.

To support Email Home, the group must have access to an SMTP server or email relay.

You need the following information:

- An email address to receive Email Home notification messages.
- An IP address and optional port for the SMTP server in the format `ip_address:port`. The default port is 25.

  You can enter up to three IP addresses. The group uses one SMTP server or email relay at any time. The first server that you specify is the default server. The group uses the other servers in the order specified, if the default server is not available.

- An email address to send (the address that appears in the *From* field in the notification email). You can use the group name at your company's email address. For example: `GroupA@company.com`

  When the intended recipient receives email, the email itself specifies which group it came from. This information is helpful in multigroup environments, and reduces the chance that the email server or recipient will discard or reject notifications.

Configure or Change Email Home Notifications

You can use Email Home to notify Dell Technical Support of hardware failures and firmware updates.

NOTE: If you want to manually send diagnostic reports to Dell Technical Support, use the Generate and Email Diagnostics wizard.

1. Click **Group** → **Group Configuration**.
2. Click the **Notifications** tab to open the Email Event Notifications panel.
3. Select the **Send email alerts to Customer Support (email home)** checkbox.
4. Enter the receiving email address in the **Local contact email** field.
5. In the **SMTP servers** section:
   a. Click **Add** to open the Add SMTP server dialog box.

      You can configure up to three servers, and you must have at least one server configured to send email.

   b. In the **IP address** field, enter the IP address for the SMTP server.

      The format is *A.B.C.D*, where A, B, C, and D are 1- to 3-digit numbers.

   c. In the **Port** field, type the port number to use to communicate with the specified server, or click the **Use default port** button.

      The default port, 25, is used if no port is defined.

   d. Repeat steps a through c to add more servers.

   e. To change the order of the listed servers, select an IP address and click the **Up** or **Down** link to reposition the address in the list.

      Servers are listed in the order that they should be used for email.
6. Click **OK** and then type an email address in the **Sender email address** field.

When you first enable Email Home, the group sends the local contact email address a confirmation message. If you do not receive this message:

- Make sure that you specified the correct information in the Email Event Notifications panel.
- Examine the PS Series event log. If no errors are logged, contact your support provider. If you have a service agreement, your support provider can help you resolve the problem.

To disable the Email Home feature, clear the **Send email alerts to Customer Support (email home)** checkbox.

Configure syslog Notification

When properly configured, the group logs events to a remote syslog server. You can then access events from the syslog server. For example, you can log events to the syslog server provided by SAN Headquarters. The syslog server must be configured to store remote log files.

NOTE: The SAN Headquarters server also includes a syslog server. You can configure groups to log events to this syslog server, including hardware alarms and performance alerts.

1. Click **Group** → **Group Configuration**.
2. Click the **Notifications** tab to open the Event Logs panel.
3. Select **Send events to syslog servers**.
4. Click **Add** under Syslog Servers.
5. Specify IP addresses for up to three syslog servers. (All the servers receive events.)
6. Under Event Priorities, select the event priorities that result in syslog server notification.
7. Click **Save all changes**.

Change the syslog Notification Configuration

1. Click **Group** → **Group Configuration**.
2. Click the **Notifications** tab to display the Event Logs panel.
3. Make any of the following changes:
   - To disable syslog notification, clear **Send events to syslog servers**.
   - To modify the IP address for a syslog server:
     1. Select the IP address and click **Modify**.
     2. Change the address and click **OK**.
   - To delete a syslog server, select the IP address and click **Delete**.
   - To change the event priorities that result in notification, select the priorities.
4. Click **Save all changes** to apply the changes.

Access the Event Log File on a Remote Computer

If you configured syslog notification, the group logs events to one or more syslog servers. The way in which you access the events depends on the syslog configuration.

Enable or Disable the Display of INFO Event Messages

You can enable (the default) or disable the display of informational messages in the Group Events window and on the CLI console. However, the group continues to log these messages.

1. Click **Group** → **Group Configuration**.
2. Click the **Notifications** tab to display the Event Logs panel.
3. Select or clear **Enable automatic delivery of INFO messages**.
4. Click **Save All Changes** to apply the changes.

5

Data Security

You can secure data at the group, volume, or NAS container level.

Table 12. Data Security Operations

Group Security	Volume Security	NAS Container Security
Enable or disable GUI and CLI access Create, modify, or delete administration accounts Add, modify, or remove a Microsoft Active Directory server Add Access Policies to the group Add Active Directory groups Enable or disable single-sign-on authentication to a group Add or modify SNMP authentication Add, modify, or delete VDS/VSS authentication Enable or disable group access options via FTP/Telnet/SSH (Advanced) Configure IPsec (Advanced) Configure a dedicated management network	Create, modify, or delete access control policies Configure, modify, or delete CHAP authentication Configure, modify, or delete iSNS servers Enable or disable iSCSI discovery filters Allow or disallow multihost access to targets	Modify UNIX permissions for Windows files or Windows directories

See the following topics for more information:

- About Group-Level Security
- About Volume-Level Security
- About NAS Container Security

6

About Group-Level Security

Group Manager supports several strategies to ensure that only the people and applications that have approved credentials can log in to the PS Series group and gain access to your data. Security can be accomplished through the following methods:

- Administration accounts — You can assign several predefined levels of administrative accounts to provide individuals with various levels of access to Group Manager’s features. To log in to the group, you must have a valid group administration account. Different account types provide different privileges. The default account, grpadmin, provides all privileges.
- RADIUS authentication — You can control access to a group and its volumes by using administration accounts to log in to the group. Use a RADIUS authentication server to enable you to centralize account management.
- Active Directory/LDAP — You can authenticate administrator sessions using LDAP. You can give group, pool, or volume administrator privileges to individual Active Directory users or to entire Active Directory groups.
- SNMP — Simple Network Management Protocol (SNMP) enables read-only access to the group. SAN Headquarters (SAN HQ) uses SNMP to retrieve data from a group.
- VDS/VSS access control — Enables Windows VDS and VSS access to the group. You must create at least one VDS/VSS access control policy that matches the access control credentials you configure on the computer by using Remote Setup Wizard or Auto-Snapshot Manager/Microsoft Edition.

To control access to data at the volume level, you can enable authentication at the iSCSI level.

Enable or Disable GUI and CLI Access

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab to open the Access panel.
3. Enable or disable the GUI or CLI access options and network services.
4. Click the Save All Changes icon in the upper-right corner of the Group Configuration window.

NOTE: The CLI provides detailed control over the group’s use of cryptographic protocols. For more information, see the grpparams crypto-legacy-protocols command in the *Dell EqualLogic Group Manager CLI Reference Guide*.

Switch Administration Authentication Type

To switch the authentication type for the group:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, select either:
   - **Local only** — Uses local authentication and local administrator users only.
   - **Active Directory** — Uses LDAP authentication, users, and groups in addition to local authentication and administrator accounts. When this option is selected, the **LDAP settings** button is available.
   - **RADIUS**— Uses RADIUS authentication and users in addition to local authentication and administrator accounts.

About Administration Accounts

Administration accounts provide various levels of access to Group Manager’s features. You must have a valid group administration account in order to log into Group Manager and gain access to a group.

If your environment requires additional security, you might consider a dedicated management network. (See Configure a Management Network for more information.)

Administration accounts allow you to specify how much control individual administrators will have over the PS Series group, according to their account type:

- Group administrators (all permissions)
- Read-only accounts (read access only to a group and can selectively enable configuration/diagnostic collection)
- Pool administrators (manage only selected pools, and if group read-only, can enable configuration/diagnostic collection)
- Volume administrators (create and manage owned volumes in selected pools)

Administration accounts can be managed locally or remotely:

- Local accounts — If you have relatively few administration accounts, this method is practical because account authentication occurs within the group. The default administration account, grpadmin, is a local account created automatically when the group is first configured.
- Remote using Active Directory (LDAP) — If you use Active Directory in your environment, you can configure a group to use LDAP to authenticate administration accounts. You can grant group, pool, or volume administrator privileges to individual Active Directory users or to entire Active Directory groups.
- Remote using a RADIUS server — If you have a large number of administration accounts, you can use an external Remote Authentication Dial-in User Service (RADIUS) server to authenticate administration accounts.

NOTE: You cannot simultaneously use RADIUS and Active Directory to authenticate administrator accounts. However, you can always add local accounts.

The default administration account, grpadmin, provides full access to Group Manager's features and allows you to perform all group operations. Some operations, such as upgrading array firmware, can be performed only by the grpadmin user.

NOTE: Dell recommends that you set up an account for each administrator, with no users sharing a single account. Further, Dell recommends that the group administrator monitor the activity of other accounts.

Types of Administration Accounts

Table 13. Types of Administration Accounts lists administration account types and their privileges. The attributes can be applied to both local accounts and Active Directory accounts or groups.

Table 13. Types of Administration Accounts

Account Type	Description
grpadmin	Can perform all group management tasks, including managing the group, storage pools, members, NAS clusters, volumes, and accounts. Group Administrator can also enable `secure erase` to securely erase data so that it cannot be recovered. Only the grpadmin account can update member firmware or fetch diagnostic files using FTP. You cannot rename, delete, or change the account type for the grpadmin account.
Group administrator	Can perform the same tasks as the grpadmin account, except updating member firmware.
Read-only	Can view information about all group objects except NAS clusters, but cannot change the group configuration. Read-only users can also save diagnostics and save the group configuration.
Pool administrator	Can view the volumes, members, snapshots, and other objects only in the pool or pools for which the account has authorization. They cannot manage members. Optionally, pool administrators can view information about all group objects except NAS clusters. Pool administrators can assign volumes to volume administrators, provided that the pool administrator has access to the pool containing the volumes, and the volume administrator has sufficient free quota space. Pool administrators cannot change the resources to which they have access.

Account Type	Description
Volume administrator	Volume administrators are (optionally) assigned a quota of storage to manage within one or more pools. They can create and manage volumes within their quota, and can perform all operations on volumes they own. Volume administrators cannot exceed their quotas by creating or modifying volumes, and cannot be assigned volumes by group or pool administrators if the capacity of the volume exceeds the free space within the quota. Volume administrators cannot modify their quotas, reassign volumes to other administrators, or change the pools or replication partners to which they have access. Volume administrators can change volumes to which they have access. Volume administrators can manage access policies and access policy groups for the volumes under their control. Volume administrators can view information only for pools and volumes to which they have access. For security purposes, the volume administrator has a limited view of group and pool configuration settings, and cannot view information, such as the SNMP community name or event log, that might enable them to gain additional access. Volume administrators also cannot view NAS clusters. Group and pool administrators can assign existing volumes to a volume administrator. If a volume is assigned to another administrator account, the volume administrator can no longer view or modify it.

Administrator accounts have the following additional restrictions:

- You cannot change the name of an administration account. Instead, you must delete the account and then recreate it with the new name.
- Only group administrator accounts can modify the pools for a pool administrator; the volume assignments, pools, quotas, or replication partners for a volume administrator; or enable or disable any account.
- Only a group administrator can modify the attributes of another group administrator account (including changing it to a read-only account), with the exceptions noted above for the default grpadmin account.
- You cannot apply read-only permission to a volume administrator or pool administrator account. Only group administrator accounts can set or remove the read-only flag.
- A pool administrator can see all volumes in their pools. The pool administrator can unassign any volume in their pools. However, the pool administrator cannot change any volume administrator’s pool access privileges or storage quotas.
- An existing account (for example, a group administrator) cannot change its type (for example, to volume administrator or pool administrator). If you need to change the privileges on an account, delete the existing account and create a new one of the desired type.

Any account can modify the following attributes of its own account:

- Contact name
- Description
- Email address
- Mobile number
- Phone number
- Password

NOTE: Active Directory accounts cannot modify their passwords through Group Manager.

Differences Between Authentication Methods

Depending on the size of your environment, the form of authorization that you choose for administrator accounts can have advantages or disadvantages. Table 14. Differences Between Authentication Methods compares various approaches.

Table 14. Differences Between Authentication Methods

Type	Advantages	Disadvantages
Active Directory groups	• Good scalability for large environments with many users; you can quickly add many administrator accounts to the group. For example, if a company hires new IT staff, and the "IT Users" group has access to the group, no extra action is required on the part of the group administrator. • Useful in environments with many PS Series groups; you can configure all groups to use the same LDAP authentication server, thus eliminating the need for maintenance of parallel sets of local accounts. • If users are removed from the Active Directory group, you do not need to update the array's list of administrator accounts to revoke access to the group.	• Active Directory administrator, not PS Series group administrator, controls which user accounts are in the group. • If the Active Directory/LDAP server is inaccessible, Active Directory accounts cannot be authenticated and logins will fail.
Active Directory or RADIUS users	• Good for smaller environments in which only a few Active Directory or RADIUS accounts are added. • PS Series group administrator controls which user accounts are in the group.	• If users are removed from the Active Directory group, the accounts remain in the PS Series group, counting against the maximum number of user accounts. • The group administrator must manually remove unused Active Directory and LDAP accounts.
Local accounts	• Good for environments with a small IT staff, or in cases where a small number of ad-hoc accounts are needed. • PS Series group administrator controls which accounts are in the group.	• Using Active Directory and RADIUS provides superior scalability to using local accounts. • Frequent changes to the roster of administrator accounts require the group administrator to make frequent updates. • If many PS Series groups are in the environment, parallel sets of administrator accounts must be created to grant administrator access to all groups.

Administration Account Attributes

Table 15. Administration Account Attributes describes the fields used in creating or modifying an administration account. You might find it beneficial to gather the information for the fields before creating an account.

The default administration account (grpadmin) is the only account capable of performing all group operations, and it is also the account you must use to perform firmware updates.

Table 15. Administration Account Attributes

Attribute	Description
Account Name	Name of the account, up to 16 alphanumeric characters. These characters are also allowed: period (.), hyphen (-), and underscore (_). The first character must be a letter or number. The last character cannot be a period. Active Directory account names can be up to 511 ASCII characters. If you enter user names containing pound signs (#) in the Group Manager CLI, the group only processes the characters preceding the pound sign; the pound sign, and all characters following it, are treated as a comment. For example, if you try to create an account named `AdminUser#Account`, the resulting account is named `AdminUser`. The Group Manager GUI does not allow you to input pound signs when creating user names.

Attribute	Description
	NOTE: Dell recommends that administrator account names not be reused after they have been deleted. All accounts can always view their own audit log information, and new accounts with previously used account names will be able to view audit records for the old account.
Password	Password for the account can be 3 to 13 ASCII characters and is case-sensitive. Punctuation characters are allowed, but spaces are not. Only the first 8 characters are used; the rest are ignored (without a message). You are not required to provide passwords for Active Directory accounts and groups after Active Directory has been configured. The passwords for these accounts are managed by the Active Directory server.
Description	Optional description for the account. Descriptions can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
Account type	Can be one of group administrator, volume administrator, pool administrator, or a read-only account.
Pool access	Pools to which the account has access and, if the account is a volume administrator, the storage quota the account can manage within the selected pools. Applies to pool administrators and volume administrators.
Additional access	Grants the pool administrator read access to the entire group. Volume administrators have read access only to the individual pools containing the storage quotas that they manage. In addition, you can grant read-only users access to collect array diagnostics and/or save configurations.
Enable account	Whether the account is active (enabled) or not. A user cannot log in to a disabled account.
Contact Information	Name, email address, and phone numbers for the account owner. Contact name can be up to 63 bytes. Email, Phone, and Mobile information can be up to 31 ASCII characters.

About Security Access Protocols

The PS Series group supports security protocols SSL/TLS and SSH, with a range of encryption algorithms. The protocols and algorithms enabled by default include some older protocols (such as SSH v1 and SSL v2) and encryption algorithms that are no longer recommended as best practices. The PS Series group supports SCP (secure copy) for copying firmware updates and diagnostic files to and from the array. It is a secure alternative to FTP and Telnet. PS Series arrays also support IPSec protocols to provide IPSec authentication and protection between group member arrays as well as between iSCSI initiators and the group. IPSec protocols must be manually enabled using the CLI. IPSec can be enabled for a group only if all members of that group support IPSec. For more information, see About IPsec.

Unless you need to enable access from older clients (web browsers or SSH clients) that do not support the current encryption protocols and authentication algorithms, Dell recommends that you disable the legacy protocols and algorithms for best security.

You must use the CLI to disable the legacy protocols; see the grpparams crypto-legacy-protocols command in the *Dell EqualLogic Group Manager CLI Reference Guide*. You can also enable or disable SSH v1 protocol support; see the grpparams cliaccess-ssh command.

SSH Key Pair Authentication

SSH key pairs are two cryptographically secure keys that can be used to authenticate a client to an SSH server. Each key pair consists of a public key and a private key. The private key is retained by the client and can be encrypted on disk with a passphrase only. The associated public key can be used to encrypt messages that only the private key can decrypt.

Limitations

SSH key pair authentication supports the following:

- RSA and SSH 2 keys
- One public key per user

- A maximum of 4096 bit.
- Minimum key length of 128 bytes.
- Local users only.

To create or view the SSH public key:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Accounts and Groups panel, select either:
    - **All accounts and groups** to view both local and remote accounts.
    - **Local accounts** to view local accounts only.
    - **Locally authenticated users** to view users that have been locally authenticated.
4. Select the account and click **Modify**. The Modify Administration Account dialog box opens.
   In the dialog box, use the **SSH Public Key** tab to view or change attributes of the public key.

   NOTE: The SSH Public Key tab is available only when all the members of the group are running firmware v9.1.x or later.
5. In the **Public key** field, enter the value for the account.
6. (Optional) In the **Description** field, enter a description for the public key.
7. Click **OK**.

Minimum Requirements for Administrative Access

To implement any form of system security and take full advantage of Group Manager's other administrative tools, you must:

- Be able to establish network access to the group through the group IP address or dedicated management address.
- Configure a group administration account with sufficient permissions to manage administrative access. The default account (grpadmin) provides this permission.
- Enable the required access options in the Group settings Administration tab for the security method being implemented.

Create a Local Administration Account

You configure, manage, and authenticate local administration accounts within the group. Local accounts are practical when you need only a small number of administration accounts for the group.

The grpadmin account is the default administration account. Dell recommends that you set up an additional account for each administrator, and reserve the default grpadmin account for maintenance operations such as firmware updates.

Accounts can be configured to be authenticated through the PS Series group or using LDAP with Active Directory.

Prerequisites

Depending on the type of account you create, you can select the following options when you create the account:

- Select one or more pools the account can manage
- Specify the quotas for each pool
- Determine whether the account has read-only access to the entire group
- Determine whether the account has access to save diagnostics and access to save config
- Select replication partners, provided you have configured replication partners

NOTE:

- Account Name, Password, and Contact information must be ASCII characters only. Description can be up to 127 Unicode characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- Dell recommends that administrator account names not be reused after they have been deleted. All accounts can always view their own audit log information, and new accounts with previously used account names will be able to view audit records for the old account.

To create a local administration account:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Accounts and Groups panel, click **Add**. The Create Administration Account wizard opens.
4. Complete the wizard steps to specify the settings for the new account and then click **Finish**.

Refer to the online help resources that are available in the user interface to view descriptions of individual fields.

Modify Local Administration Accounts

This procedure applies only to modifying local administration accounts and has the following constraints:

- You cannot change the account name. Instead, you must delete the account and then recreate it with a new name.
- You cannot disable, delete, change the name, or change the type of the default group administration (grpadmin) account.
- Except in one case, you cannot change the account type; instead, you must delete the account and recreate it with the new account type. The exception is that you can change a grpadmin account (other than the default) to read-only and a read-only account to grpadmin.

To change a local account:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab to display the Accounts and Groups panel.
3. In the View drop-down list, select **Local Accounts**.
4. Select the account and click **Modify**. The Modify Administration Account dialog box opens.
   - To change the account password or description, click the **General** tab and change the General Settings information.
   - To change the account contact information, click the **Contact** tab and change the information.
   - To make changes to the account type, pool access, additional access, or to disable the account, click the **Permissions** tab.

     This tab is not available to grpadmin accounts because group administrators must always have full access.
   - To change or select the partners that can be used for replication, click the **Replication Partners** tab.

     This tab is available to volume administrators only.
5. Click **OK**.

Delete Local Administration Accounts

You can delete an administration account when it is no longer needed.

NOTE: If you delete a volume administrator account, the volumes it manages are not deleted, and its replication and operations continue as scheduled.

To delete an administration account or group:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the **Accounts and Groups** panel, select the account or group you want to delete.

4. Click **Delete** and confirm that you want to delete the account.

NOTE: Dell recommends that administrator account names not be reused after they have been deleted. All accounts can always view their own audit log information, and new accounts with previously used account names will be able to view audit records for the old account.

About RADIUS Accounts

If you have a large number of accounts, you can use an external RADIUS server to simplify account management and centralize the management of administration accounts. RADIUS is the abbreviation for Remote Authentication Dial-In User Service, which provides a central authorization and authentication service for all access requests.

The RADIUS server authenticates administration accounts and also determines the account privileges. You can also use a RADIUS accounting server to monitor the login and logout times for accounts that a RADIUS server authenticates.

If you use Active Directory in your environment, you can also configure the group to use LDAP to authenticate administration accounts. You can grant group, pool, volume administrator, or read-only privileges to individual Active Directory users or to entire Active Directory groups.

RADIUS servers are implemented in different ways. Depending on your implementation, a RADIUS server can verify account credentials against a local database, or it can verify them against an external resource such as a Microsoft Windows Active Directory service domain.

NOTE: External administration accounts depend on the availability of the RADIUS server and any related resources. If these resources are not available, accounts cannot be authenticated and a login does not succeed.

For information about using Active Directory to manage and authenticate administration accounts, see the Technical Report *Using Active Directory for Account Authentication to a PS Series Group* on the customer support website.

For other RADIUS implementations, see your RADIUS server documentation for information about setting up the RADIUS server and configuring vendor-specific attributes (VSAs).

You can use multiple RADIUS authentication servers for increased availability.

RADIUS Attributes for Administration Accounts

A RADIUS server uses attributes to authorize accounts as group administrator, pool administrator, volume administrator, or read-only, and to store account contact information.

For security reasons, Dell recommends that you require vendor-specific attributes. See your RADIUS server documentation for information on how to set attributes.

For each account, you must set the `Service-Type` attribute to one of these values:

- `EQL-Admin-Privilege`—Specifies that the account is either a group administrator account, a pool administrator account, or a volume administrator account. If you do not specify the `EQL-Admin-Privilege` attribute, the account defaults to group administrator.
- `NAS-Prompt`—Specifies that the account is a read-only account.

In addition, you must configure vendor-specific attributes (VSAs) for each account if you meet one of these conditions:

- You want to create a pool administrator account. You must specify the `EQL-Admin-Privilege` attribute and the `Admin-Pool-Access` attribute.
- You want to create a volume administrator account. You must specify the `EQL-Admin-Privilege` attribute, the `Admin-Pool-Access` attribute, and (optionally) the `Admin-Repl-Site-Access` attribute.

  NOTE: A replication quota must be included inside the `Admin-Repl-Site-Access` attribute for authentication to work properly.
- You want to create a read-only account. You must specify the `EQL-Admin-Privilege` attribute and the `Admin-Account-Type` attribute.

- You plan to select the `Require vendor-specific RADIUS attribute` option when you configure the group to use a RADIUS authentication server. You must specify the `EQL-Admin-Privilege` attribute.

Table 16. Vendor-Specific Attributes describes the Dell vendor-specific attributes and values for RADIUS attributes.

Table 16. Vendor-Specific Attributes

Attribute	Field	Required Value

`EQL-Admin-Privilege`
Specifies that the account is a group administrator account or a pool administrator account.
The RADIUS server must return the value of this attribute to the group in the Access-Accept message.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		6
		Decimal (`0` for group administrator; `1` for pool administrator; `2` for pool administrator with read access to the entire group; `3` for volume administrator).
		To create a read-only account, set the `EQL-Admin-Privilege` attribute to `0` and the `Admin-Account-Type` attribute to `RO`.

`Admin-Pool-Access`
Specifies the pools to which the pool administrator account has access and, for volume administrators, the account's storage within that pool.
Required if the value of the `EQL-Admin-Privilege` attribute is `1` (pool administrator account) or `3` (volume administrator account).
The quota for volume administration accounts is expressed as *`PoolNameQuota`*, with gb and mb (representing GB and MB, respectively) appended to the quota.
For example: `Pool1 25gb` sets the quota for Pool1 to 25GB, and `Pool1 500mb` sets a quota of 500MB. Use `unlimited` to set an unlimited quota for the pool (for example, `Pool1 unlimited`). If no unit is specified, the default capacity unit is MB.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		7
		String (comma-separated list of pools; 3 to 247 ASCII characters)

`Admin-Repl-Site-Access`
Specifies the sites to which the volume administrator can replicate volumes. Required if the value of the `EQL-Admin-Privilege` attribute is 3 (volume administrator account). Used only for volume administrators.
NOTE: A replication quota must be included inside the `Admin-Repl-Site-Access` attribute for authentication to work properly.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		8
		String (comma-separated list of sites; 3 to 249 ASCII characters)

`Admin-Account-Type`
Specifies whether the account is read-only (`RO`) or read-write (`RW`).
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		9
		RO or RW

`Admin-Full-Name`
(Optional) Name of the administrator using the account.
	`VSA vendor ID`
	`VSA number`
		12740
		1

Attribute	Field	Required Value
	`VSA syntax`	String (3 to 247 ASCII characters)

`Admin-Email`
(Optional) Email address of the administrator.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		2
		String (3 to 247 ASCII characters)

`Admin-Phone`
(Optional) Phone number for the administrator.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		3
		String (3 to 247 ASCII characters)

`Admin-Mobile`
(Optional) Mobile phone number for the administrator.
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		4
		String (3 to 247 ASCII characters)

`Admin-Poll-Interval`
Frequency, in seconds, the GUI polls the group configuration data. The default is 30 (seconds).
	`VSA vendor ID`
	`VSA number`
	`VSA syntax`
		12740
		5
		Integer (up to 6 numerals)

Prerequisites for Configuring RADIUS Servers

Before you use a RADIUS server to authenticate administration accounts (or CHAP accounts for iSCSI access), you must install the server and set up the accounts:

1. Install and configure the RADIUS authentication server.

   For example, to add the group as a RADIUS client on a Microsoft Windows server, you must specify the following items:

   - Name (also called Friendly Name) for the client. Dell recommends using the group name.
   - Group IP address (also called Client address) or dedicated management network IP address.
   - Vendor Name attribute. Select RADIUS Standard.
   - Password (also called Shared Secret) of up to 63 ASCII characters. This password should also be entered in Group Manager when you configure the group to use the RADIUS authentication server.

   NOTE: Though using a password is not required, Dell recommends that you use one for increased security.
2. For iSCSI CHAP (Challenge Handshake Authentication Protocol) accounts, add each configured network interface on all the group members as a RADIUS client. Specify the network interface IP address and, optionally, a password (or secret), up to 63 ASCII characters. If you specify a password, enter this password when you configure the group to use the RADIUS authentication server. Dell recommends that you use a password for increased security.
3. For administration accounts, set up the attributes that allow the server to authorize accounts as group administrator, pool administrator, or read-only accounts.
4. Set up the accounts. You can set up accounts on the RADIUS server or a different resource, such as Active Directory. The RADIUS server verifies login credentials (account name and password) that the user supplies against these accounts.

The RADIUS server must be accessible to all the group members.

Configure a RADIUS Server

When configuring a RADIUS server in a pure IPv6 environment, you must:

- Specify, on the RADIUS server, a RADIUS client for every IPv6 IP as an entry for the Microsoft Windows server.
- Enable access for the RADIUS user.

To configure the RADIUS server:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, under Authentication Type, select **RADIUS** and then click the **RADIUS settings** button to open the RADIUS Settings dialog box.
4. In the RADIUS Authentication Servers section, click **Add**. The Add RADIUS Authentication Server dialog box opens.
5. Specify the IP address of the server.

   If the server uses a port other than port 1812 (the default), specify the correct port number.
6. Specify and confirm the RADIUS secret and click **OK**.
7. (Optional) Specify the server timeout and retry values:

   - **Request timeout (seconds)** — Number of seconds the group waits for an accounting server to transmit before timing out. The default setting is 2 seconds.
   - **Retries** — Number of times the group tries to contact an accounting server after the first failure. The default setting is 1.
8. (Optional) If you want to add information for RADIUS accounting servers, click **Add** under the Radius Accounting Servers subpanel and repeat steps 5 to 7.
9. To add additional servers, repeat steps 3 to 6. To finish, click **OK**.
10. On the **Group Configuration iSCSI** tab, you must select the checkbox for enabling RADIUS authentication for iSCSI initiators.

By default, all RADIUS authentication and accounting options are now enabled.

Modify RADIUS Server Settings

You can modify the following settings on a RADIUS authentication or accounting server:

- Server IP address
- Password (secret) — up to 63 ASCII characters
- Request timeout value
- Number of retries value

To modify a RADIUS setting:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, select **RADIUS** and then click the **RADIUS settings** button to open the RADIUS Settings dialog box.
4. Select the server IP address and click **Modify**. The Modify RADIUS Authentication Server dialog box opens.
5. Change a server IP address or password settings and click **OK**.
6. Click **OK** to confirm.

Delete a RADIUS Server Connection

To delete an unwanted RADIUS server connection:

NOTE: To delete a RADIUS account, remove it from the Active Directory server first before you delete it from the storage group. Deleting an Active Directory account (or group) from the storage group does not remove it from the Active Directory server.

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, select **RADIUS** and then click the **RADIUS settings** button to open the RADIUS Settings dialog box.
4. Select the server IP address that you want to remove and then click **Delete**.
5. To confirm the deletion, click **Yes**.
6. Click **OK**.

About LDAP Authorization and Active Directory

LDAP is the abbreviation for Lightweight Directory Access Protocol, which provides a simplified protocol for authenticating users. An LDAP server typically contains a database of users, user names, passwords, and related information. LDAP clients are able to interrogate the server to authenticate these users and obtain the account characteristics.

Active Directory is an LDAP-compliant database that contains objects (typically users, computers, and groups) and provides authentication and authorization mechanisms in which other related services can be deployed.

If your environment uses Active Directory, you can authenticate administrator sessions using LDAP. Individual Active Directory users, or entire Active Directory groups, can be given group, pool, or volume administrator privileges.

To use LDAP authentication, you must first configure the group to communicate with one or more LDAP servers.

The Active Directory Configuration wizard enables you to configure NTP and DNS or modify the existing NTP or DNS configuration. You can also perform these tasks at a later time. See the*Dell EqualLogic PS Series Storage Arrays Release Notes* for more information about NTP requirements for using Active Directory in a NAS cluster.

To use Active Directory in a NAS cluster:

- The Active Directory server and the PS Series group must use a common source of time.
- You must configure the NAS cluster to use DNS. The DNS servers you specify must be the same DNS servers that your Active Directory domain controllers use.

Add an Active Directory Server

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, select **Active Directory** as the authentication type.
   - If no Active Directory servers have been added yet, the Active Directory settings dialog box opens.
   - If one or more Active Directory servers have already been added, click **AD settings** to open the Active Directory settings dialog box.
4. In the Active Directory settings dialog box, click **Add**. The Add List Item dialog box opens and prompts you to enter the AD server's IP address.
5. Type in the IP address for the Active Directory server and click **OK**. The IP address appears in the list of Active Directory servers.

Configure Active Directory Authentication

To configure LDAP authentication for the group:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, set the authentication type to **Active Directory** and click **AD settings** to display the Active Directory Settings dialog box.
4. In the Active Directory servers section, click **Add**. The Add List Item dialog box opens.
5. Type the IP address of the Active Directory server and click **OK**.
6. Repeat steps 3 to 5 to add up to three IP addresses.

   NOTE: Adding multiple Active Directory servers ensures continued authentication of Active Directory accounts even in the event of a resource outage. The group uses the first Active Directory server in the list for authenticating accounts; if the group cannot establish contact with the first server, it uses the other Active Directory servers to authenticate administrator logins.
7. Select the Active Directory server that you want to configure.
8. In the AD server settings section, select **Secure protocol:** and choose **TLS or none**.

9. Select whether to use the default port for the selected protocol, or specify a different port.
10. Type the Base DN for the Active Directory server, or select **Get Default** to use the default value. The Base DN can be up to 254 ASCII characters.
11. Select whether to use anonymous connections to the server or type a Bind DN.
12. If a Bind DN is specified, type the Bind password. Passwords can be up to 63 ASCII characters.
13. To test the new Active Directory settings, click the **Test AD settings** button. Group Manager tests the Active Directory settings for all servers. If authentication fails, a dialog box opens, listing the Active Directory servers with which connections could not be established. If no connections can be established, you can accept the configuration as is or click **Cancel** and check the Active Directory settings again.
14. Click **OK**.

Modify Active Directory Accounts and Groups

When you modify Active Directory accounts and groups, the following restrictions apply:

- You cannot change the account name. Instead, you must delete the account and then add it back with the updated name in Active Directory.
- You cannot modify cached accounts. You can only view their configuration details.
- You cannot change the account type. Instead, you must delete the account and recreate it with the desired account type.

When you modify Active Directory groups, the following considerations apply:

- An Active Directory security/distribution group is added to the PS Series group with the attribute that all members of the AD group now have access. If changes are made to any members of the group, the changes are automatically integrated the next time the members log in to the group.
  - When a new user is added to the Active Directory group, the user automatically has access to the group.
  - When an Active Directory user is removed from the AD group, the user no longer has access to the group.
  - When the user name of a current member of the AD group is modified in Active Directory, no changes need to be made for that user on the PS Series group.
- When you change the name of the Active Directory group, the group must be deleted from the PS Series group and then re-added with the new name.

To change an Active Directory account or group:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Accounts and Groups panel, select either:
   - **All accounts and groups** to view both local and remote accounts.
   - **Active Directory users** to view only Active Directory user accounts.
   - **Active Directory groups** to view only Active Directory group accounts.
4. Select the account and click **Modify**. The Modify Administration Account dialog box opens.

   In the dialog box, use the Account type section to change attributes of the account type:

   - If the account type is Pool administrator or Volume administrator, you can use the Pool access section to specify the pools to which the account has access and the storage quota for the account.
   - If the account type is Pool administrator, you can use the Additional access section to give the account read-only access to the entire group.

   You can also grant read-only accounts permission to save diagnostics and save config from this dialog box.
5. To change replication partners for a volume administrator, click the **Replication Partners** tab and change the selections.

   NOTE: Only users with group administrator privileges can modify the NAS container replication configuration.
6. Click **OK**.

Test the Active Directory Server

After you have added the Active Directory server, test your connection by clicking **Test AD settings**. The firmware tests all of the Active Directory servers in the list and reports the results of each connection attempt.

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. In the Access panel, make sure that the **Enable web access** checkbox is selected and select **Active Directory** as the authentication type.
4. Click **AD settings** to open the Active Directory Settings dialog box.
5. Select the IP address of the server you want to test.
6. Confirm that the AD server is correctly configured and click the **Test AD settings** button.
   - A successful connection shows the IP address of the server and the messages `Connection established/Test Search succeeded`.
   - If the connection is unsuccessful, the dialog box reports a`Failed to Connect` error.
7. Click **OK**.

Set Automatic Login Preferences

To set automatic login preferences:

NOTE:

- Make sure that the PS Series group is configured for single sign-on using Active Directory before you log in using your Active Directory credentials.
- When using single sign-on, the group name cannot contain more than 19 characters.

1. Click **Tools** to open the drop-down menu and select **User preferences**.
2. In the Modify User Preferences dialog box, click the **Communication** tab.
3. In the Connection Policies panel, select the **Automatically log in using Windows Active Directory credentials** checkbox.
4. Click **OK**.

You can also select the **Automatically reconnect if disconnected** and **Keep session alive when temporarily leaving GUI page** options.

Promote or Demote the Active Directory Server

To provide for fault tolerance, an administrator can configure the PS Series group with the IP addresses of up to three Active Directory (AD) peers. When a user logs in to AD, the PS Series group attempts to contact the AD servers in the order in which they are listed. If an attempt to connect to the first server fails, the PS Series group tries to connect to the second, and so on.

An administrator can adjust the order in which AD servers appear by promoting or demoting an AD server in the server list.

To promote or demote a server to a higher position in the list:

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. Click **AD settings** button to display the Active Directory settings dialog box.
4. Select the IP address of the server in the list of Active Directory servers, and then either click **Up** to promote the server or click **Down** to demote the server.

Remove an Active Directory Server

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, select **Active Directory** for the authentication type and then click **AD settings** to open the Active Directory settings dialog box.
4. In the Active Directory servers panel, select the AD server that you want to remove and click **Delete**.
5. Click **OK**.

About Active Directory Groups

In addition to local and RADIUS administration, administrator account sessions can be authenticated using Active Directory. Individual Active Directory users, or entire Active Directory groups, can be given access to Group Manager using the same levels of access permission available for local user accounts.

Using Active Directory authentication is useful in large SAN environments in which administrators require access to multiple groups. By configuring each PS Series group to use the Active Directory server, you do not need to maintain parallel sets of local accounts for each group.

You can configure the group to authenticate accounts using multiple Active Directory servers; if the primary Active Directory server is unavailable due to a connection issue, outage, or disaster event, the extra servers will ensure continued Active Directory authentication of administrator accounts.

You can also use Active Directory authentication as an alternative to RADIUS authentication.

To use Active Directory authentication, you must first set the group's authentication type to Active Directory, and add one or more Active Directory servers. If you are using Active Directory for authentication, you cannot use RADIUS authentication for the group. You can, however, still create and use locally authenticated user accounts.

Add Active Directory Groups

To add all accounts in an Active Directory group to the list of administrator accounts:

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Accounts and Groups panel, click **Add**. The Create administration account dialog box opens.
4. Complete the wizard steps to specify the settings for the new account and click **Finish**.

Add an Active Directory User to the Group

Before an Active Directory (AD) user can use single sign-on (SSO) to automatically log in to the PS Series group, an administrator must grant that user permission. You perform the same procedure to grant access to AD groups.

1. Click **Group → Group Configuration**.
2. Click the **Administration** tab.
3. In the Accounts and Groups panel, click **Add**. The Create Administration Account wizard opens, showing step 1 – General Settings.
4. Select **Active Directory user**. To add an AD group, select **Active Directory group**. (When using single sign-on, the group name cannot contain more than 19 characters.)
5. In the General Settings section, specify the user name of the PS Series group for the AD user. Each user name must be unique. Click the **Check name** button to make sure the name that you specified is not already in use.
6. Complete the required fields in each remaining step of the wizard until you reach the Summary page.
7. Confirm that the settings are correct and click **Finish** to create the AD user.

NOTE: If you log in to Windows using your Active Directory credentials, you will be logged in to the PS Series group automatically without re-authenticating.

Change the Active Directory Group Name

Before you change the name of a PS Series group that has already been configured for single sign-on, Dell recommends that you leave the current Active Directory (AD) domain, change the group name, and then join the AD domain again using the new name.

NOTE: When using single sign-on, the group name cannot contain more than 19 characters.

About Single Sign-On

Single sign-on enables users who have already logged in to their PCs using Windows Active Directory credentials to automatically log in to the Group Manager GUI without having to specify the Windows Active Directory login credentials again. To use single sign-on, configure the PS Series group to direct it to the same Active Directory domain that authenticates users when they log in to their workstations.

NOTE:

- To use this option, you must specify an Active Directory server on the network.
- When using single sign-on, the group name cannot contain more than 19 characters.

Check Server and Single Sign-On Configuration

After you have successfully configured and tested the Active Directory server and single sign-on, you can check the status of these features in the Group Manager GUI.

To check the Active Directory server and single sign-on connections:

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, both AD Server and Single Sign-On should display the status as `Configured` when they have been set up correctly. To verify that each feature is connected, click **Check**.
   - The AD Server **Check** button queries the configured IP address and port and returns the results in a dialog box.
   - The Single Sign-On **Check** button connects with Active Directory and verifies that the PS Series group has joined the Active Directory domain that you configured.

Use Single Sign-On to Log In to a PS Series Group

Before you can use single sign-on (SSO) and Windows Active Directory (AD) to log in to the PS Series group with your Windows AD credentials, make sure that the group has been configured for single sign-on.

1. Configure the PS Series group for single sign-on and join it to an Active Directory domain (for example, the Engineering domain).

   NOTE: For single sign-on, configure the PS Series group to direct it to the same Active Directory domain that authenticates users when they log in to their workstations.
2. Log in to a client computer that has already joined the AD domain (for example, Engineering), using your user credentials from Active Directory.
3. Start the Group Manager GUI. If you are logging in for the first time using AD and you have administrator permissions, the **Use Windows Credentials** option appears in the Log In to Group dialog box.

   NOTE: If Use Windows Credentials does not appear in the dialog box, make sure that the PS Series group has been configured for single sign-on, and that the group is in the same AD domain as the client computer that you are using to log in.
4. Select **Use Windows Credentials** and click **OK**.
5. If login is successful, the GUI prompts whether to automatically use single sign-on credentials for all future logins.
6. Select **Yes** to automatically log in with single sign-on, or **No** to retain the option to select the login method for the next session. The GUI directs you to the Group Manager home page.

After you are logged in, your AD user name appears at the top of the Group Manager window.

For subsequent login sessions using SSO, you will be logged in automatically as soon as you start the Group Manager GUI.

Log Out of Group Manager Using Single Sign-On

Logging out of the Group Manager GUI using single sign-on (SSO) is the same as logging out using regular login credentials, but you see additional options.

1. Click **Log out** at the top-right corner of the Group Manager GUI. A dialog box opens.
2. Select **Log out from Group Manager GUI** to log out of Group Manager, or select **Log in as different user** to log in using your regular login credentials.

Enable or Disable Single Sign-On

To enable single sign-on:

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, under Active Directory Configuration, click the **Join domain/Enable SSO** button.
4. In the dialog box that opens, type the Active Directory administrator user name and password to join the Active Directory domain.
5. Click **Join**.

To disable single sign-on:

1. Click **Group** → **Group Configuration**.
2. Click the **Administration** tab.
3. In the Authentication panel, under Active Directory Configuration, click the **Leave domain/Disable SSO** button.
4. In the dialog box that opens, type the Active Directory administrator user name and password to leave the Active Directory domain.
5. Click **Leave**.

About SNMP Access to the Group

You can use Simple Network Management Protocol (SNMP) for read-only access to a PS Series group through one or more read-only community names.

NOTE: SAN Headquarters requires you to configure SNMP access to a group. The Manual Transfer Utility (MTU) requires you to configure SNMP access and specify `public` for the SNMP community name. In addition, SAN Headquarters uses SNMP to retrieve data from the group to display status.

About SNMP Authentication

Some PS Series group operations, such as sending traps about hardware issues, require you to configure SNMP authentication between the group and an SNMP server.

Add or Change SNMP Access to a Group

To add an SNMP community name:

1. Click **Group** → **Group Configuration**.
2. Click the **SNMP** tab.
3. Click **Add** in the SNMP Access panel.
4. Enter an SNMP community name (for example, `public`).

   You can specify up to 5 names, and each name can be up to 64 ASCII characters long. Names cannot contain the following characters: space, tab, comma, pound sign (#).
5. Click **OK**.

6. Click **Save all changes**.

To change or delete an SNMP community name:

1. In the SNMP Access panel, select the name.
2. Click **Modify** or **Delete** as needed.

   You can specify up to 5 names, and each name can be up to 64 ASCII characters long. Names cannot contain the following characters: space, tab, comma, pound sign (#).
3. Click **OK**.
4. Click **Save all changes**.

Display SNMP Access to a Group

1. Click **Group** → **Group Configuration**.
2. Click the **SNMP** tab.

About SNMP Traps

SNMP traps are unsolicited event messages sent to a management console by an agent. PS Series arrays send traps for equipment issues and security issues.

The PS Series array Management Information Bases (MIBs) contain information about SNMP traps and trap thresholds.

SNMP Traps

PS Series traps are unsolicited event messages sent to a management console by an agent. PS Series arrays send traps for equipment issues and security issues. The PS Series array MIBs (Management Information Bases) contain information about SNMP traps and trap thresholds. Table 17. PS Series SNMP Traps describes these traps.

Table 17. PS Series SNMP Traps

Trap Type	Trap Names
Battery backup	`eqlMemberHealthBatteryLessThan72Hours`, `eqlMemberHealthNVRAMBatteryFailed`, `eqlMemberHealthhighBatteryTemperature`
eqlgroupAdminLoginStatus	`Login-stat-success`, `login-stat-logout`, `login-stat-failure`

Component	`eqlMemberHealthhwComponentFailedCrit`, `eqlMemberHealthincompatControlModule`
	`eqlMemberHealthopsPanelFailure`, `eqlMemberHealthemmLinkFailure`, `eqlDiskStatusChange`

Fan and PSU	`eqlMemberHealthFanSpeedHighThreshold`, `eqlMemberHealthFanSpeedLowThreshold`, `eqlMemberHealthFanTrayRemoved`, `eqlMemberHealthBothFanTraysRemoved`, `eqlMemberHealthPowerSupplyFailure`
GUI/WebUI login success/failure, logout	`Login-stat-success`, `login-stat-logout`, `login-stat-failure`
iSCSI	`iscsiTgtLoginFailure`, `iscsiIntrLoginFailure`, `iscsiInstSessionFailure`, `scsiTgtDevicesStatusChanged`, `scsiLuStatusChanged`
Link	`linkUp`, `linkDown`
Member Offline	`eqlMemberStatusChanged`
eqliscsiVolumeMultiInitiatorAttributeChanged	`allowed` or `not-allowed`

Trap Type	Trap Names
Network Configuration	· MemberGatewayIPAddrChanged · NetmaskChange · eqlgroupIPv4AddrChanged · eqlgroupIPv6AddrChanged · eqlMWKAChangeNotification
RAID	eqlMemberHealthRAIDSetDoubleFaulted, eqlMemberHealthRAIDLostCache, eqlMemberHealthRAIDSetLostBlkTableFull, eqlMemberHealthRaidOrphanCache, eqlMemberHealthRaidMultipleRaidSets
Security	authenticationFailure
Start	coldStart, warmStart
Temperature	eqlMemberHealthTempSensorHighThreshold, eqlMemberHealthTempSensorLowThreshold, eqlMemberHealthlowAmbientTemp
Volume Offline	eqliscsiVolumeOfflineStatusAlert

Add SNMP Trap Destinations

To configure network addresses to receive SNMP traps from the group, you need to define the SNMP trap destination.

1. Click **Group** → **Group Configuration**.
2. Click the **SNMP** tab.
3. Click **Add** in the SNMP Traps panel.
4. Type the IP address where SNMP traps are sent and click **OK**.
5. (Optional) Modify the SNMP trap community name. The default is **SNMP-trap**.
6. Click **Save all changes**.

Change the SNMP Trap Configuration

1. Click **Group** → **Group Configuration**.
2. Click the **SNMP** tab to open the SNMP Traps panel.
3. Select the IP address and either:
   - Click **Delete** to delete an SNMP trap destination.
   - Click **Modify** → **OK** to change the IP address.
4. To modify the SNMP trap community name, change the name in the **SNMP trap community name** field.

   NOTE: Community names cannot contain the following characters: space, tab, comma, pound sign (#).
5. Click **Save all changes** to apply the modification.

About VDS and VSS Authentication

Microsoft provides the following services for storage management on Microsoft Windows computers:

- Virtual Disk Service (VDS) provides an end-to-end solution for managing storage hardware and disks, and for creating volumes on those disks. VDS provides an interface using a set of APIs that allows you to manage disks.
- Volume Shadow Copy Service (VSS) is a Windows service for capturing and creating snapshots called shadow copies. VSS operates at the block level of the file system and provides a backup infrastructure for Microsoft operating systems.

NOTE: Built-in services might need to be enabled before you can use them.

To be able to perform management operations using VDS and VSS, you must first allow these services to access your PS series group. You use the same access control methods (access policies, access policy groups, and basic access points) to define VDS/VSS access.

Display and Configure Windows Service Access to a Group

To be able to perform management operations using VDS and VSS, you must first allow these services to access your PS series group. You use the same access control methods (access policies, access policy groups, and basic access points) to define VDS/VSS access.

To allow VDS and VSS access to the group, you must create at least one VDS/VSS access control policy that matches the access control credentials you configure on the computer by using Remote Setup Wizard or Auto-Snapshot Manager/Microsoft Edition.

NOTE: Auto-Snapshot Manager/Microsoft Edition requires you to configure Windows service access to a group.

VDS/VSS access control policies use the same criteria for restricting access as iSCSI target access control policies: CHAP user name, iSCSI initiator name, or iSCSI initiator IP address.

1. Click **Group** → **Group Configuration**.
2. Click the **VDS/VSS** tab.

Add a VDS/VSS Access Control Policy

To configure VDS/VSS access, add a new access control policy:

1. Click **Group** → **Group Configuration**.
2. Click the **VDS/VSS** tab.
3. In the VDS/VSS Access Control List panel, click the appropriate **Add** button to add an access policy group, an access policy, or a basic access point:
   - In the Add Access Policy or Add Access Policy Group dialog box, select the checkbox next to the policy names you want to add, or click **New** to create a new access policy or policy group.
   - In the New Basic Access Point dialog box, enter a new access point (any combination of CHAP account name, iSCSI initiator name, or IP address), then click **OK**.
4. Click **OK**.

Modify a VDS/VSS Access Control Policy

1. Click **Group** → **Group Configuration**.
2. Click the **VDS/VSS** tab.
3. To modify a policy, select the policy from the appropriate subpanel (Access policy groups, Access policies, or Basic access points) and click the corresponding **Modify** command:
   - In the Modify Access Policy Group dialog box, change the name or description of the policy group. You can also select an individual policy for the group and click **Modify**.
   - In the Modify Access Policy dialog box, change the name or description of the policy. You can also select an individual access point for the policy and click **Modify**.
   - In the Modify Basic Access Point dialog box, change the CHAP user name, iSCSI initiator name, or IP address, then click **OK**.
4. Click **OK**.

Delete a VDS/VSS Access Control Policy

1. Click **Group** → **Group Configuration**.
2. Click the **VDS/VSS** tab.
3. Select the policy from the appropriate subpanel (Access policy groups, Access policies, or Basic access points) and click the corresponding **Delete** command.

4. Confirm that you want to delete the policy.

When you delete or modify a basic access point, you might need to update any computer that was previously accessing volumes through that access point.

About IPsec

IPsec is a set of standardized protocols designed to allow systems on IP-based networks to verify each other's identities and create secured communication links. IPsec uses cryptographic security mechanisms for authentication and protection. IPsec validates the identity of devices communicating over IP-based networks, encrypts all data passing between participating systems, and protects against disclosure, modification, eavesdropping, and attack. IPsec is supported for both IPv4 and IPv6 networks.

NOTE: For more general information about IPsec, refer to the website of the Internet Engineering Task Force (ietf.org), the organization that originally developed the IPsec protocols.

In the context of an iSCSI SAN that uses EqualLogic PS Series storage arrays, IPsec secures communications between group member arrays and also between iSCSI initiators and the group. You can use policies to configure your IPsec implementation to protect iSCSI traffic based on initiator IP address, initiators in a specific subnet, or network protocol. IPsec authentication is handled using certificates or pre-shared keys.

Types of Protected Traffic

The types of traffic protected by IPsec are shown in Figure 9. What IPsec Protects.

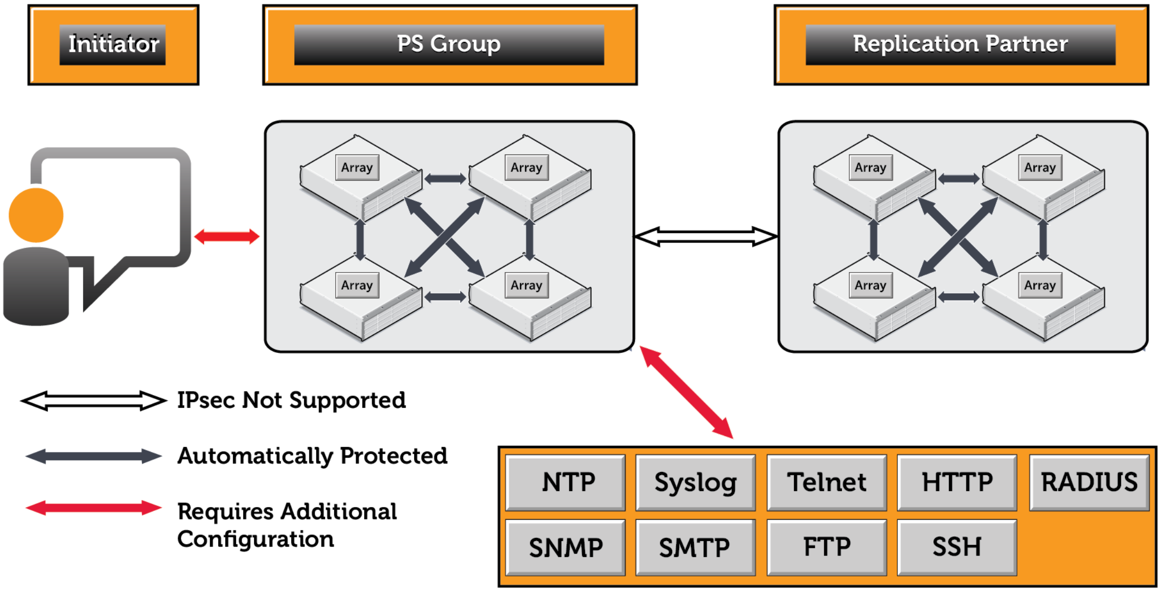

Figure 9. What IPsec Protects

IP Traffic Protection

To enable IPsec protection for traffic between the group and iSCSI initiators, use the following basic process:

NOTE: This process is not required for protecting communications between group members. After IPsec is enabled, all network traffic between group members is automatically protected, without need for further configuration.

1. A group administrator creates security parameters to specify how traffic should be authenticated.
2. A group administrator creates policies to identify traffic and determine what action to take for it:
   - Traffic is dropped.
   - Traffic is allowed to pass directly through to the array in the clear.

- Traffic is protected using certificates or pre-shared keys.

NOTE: IPsec configurations cannot be modified. They must be removed and then recreated using the new configuration.

Protect Communication Between Group Members

To enable IPsec security for communication between group members, use the ipsec enable CLI command.

After IPsec is enabled, all network traffic between group members is protected automatically. No further configuration is required.

Any incoming or outgoing IP traffic that travels between hosts and the group can be protected with IPsec. This traffic includes, but is not necessarily limited to:

- iSCSI traffic
- Telnet and SSH connections to the Group Manager CLI
- HTTP connections to the Group Manager GUI
- SMTP email notifications
- Syslog
- NTP
- RADIUS
- SNMP

NOTE: If IPsec is enabled but no security parameters or policies are in place, intragroup traffic is protected, and traffic to and from the group is allowed to pass without being protected or dropped.

IPsec During Replication

The PS Series firmware provides no mechanism for using IPsec to protect traffic between replication partners. It is technically possible to create IPsec polices on both the primary and secondary group in which each group treats the other as an iSCSI initiator and traffic is protected accordingly. However, this configuration is not supported, and Dell recommends against implementing it in a production environment.

IPsec Policies

Traffic that meets the conditions stipulated by the policy can either be passed, dropped, or protected using an IPsec security parameter associated with the policy.

You can use IPsec policies to apply IPsec protection to traffic that meets one or more of the following criteria:

- Data traveling to or from specific IP addresses, or a range of IP addresses defined by a specific subnet or netmask
- IPv4 or IPv6 traffic
- Specific network protocols: TCP, UDP or ICMP (either IPv4 or IPv6)

Unless explicitly specified by the policy, traffic is allowed to pass. If you want to drop all traffic that is not explicitly protected or passed, you must create an IPsec policy that drops traffic by default.

If multiple IPsec policies are in place, the system determines their priority by the order in which they were created. Policies created first take precedence over policies created later.

You can also use IPsec policies to determine what traffic is being protected using IPsec, and what traffic is being passed or dropped without encryption.

Security Certificates

Security certificates are used in an IPsec configuration as one method of authenticating secured connections between iSCSI initiators and the group. Implementation of an IPsec-enabled SAN requires both a root-CA (Certificate Authority) certificate from the issuing authority and a local certificate to authenticate the group.

You can generate certificates suitable for use in IPsec connections to the PS Series using any Windows, OpenSSL, or other commercial Certificate Authority product.

From the Group Manager CLI, you can import, display, and delete certificates, using the ipsec certificate commands. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

IPsec Security Parameters

IPsec security parameters control the authentication and key negotiation carried out using the Internet Key Exchange IKEv1 or IKEv2 protocol.

Security parameters specify the following features:

- Using IKEv1, IKEv2, or manual keying

  NOTE: While it is possible to configure IPsec to use manual keys via the CLI command, Dell strongly cautions that you do not use this command. Using the command can lead to extremely serious security risks. Do not use this command. Consequently, Dell strongly discourages the use of manual keying in any production environment. IKEv1 or IKEv2 are the preferred keying methods.
- Using certificates and pre-shared keys (PSK)
- Establishing Transport Mode or Tunnel Mode connections

  NOTE: Unless specifically configured, IKEv1 and Transport Mode are used by default.

IPsec security parameters are managed using the ipsec security-params commands. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

IPsec Security Associations (SA)

The pairing of an IPsec security parameter with an IPsec policy forms an IPsec security association (SA), which formalizes the secured connection between the group and a host connected to it. Each protected connection to the group is a unique security association, and each system can have multiple security associations, allowing it to have authenticated communications with many other systems.

NOTE: You can view or delete security associations using the ipsec security-association commands. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

IPsec Pre-Shared Keys (PSKs)

In addition to using certificates, you can use pre-shared keys to authenticate secured connections. Pre-shared keys are identical strings that are specified at both ends of the communications pathway. The keys enable the systems to correctly identify each other.

You can use either ASCII or hexadecimal strings. ASCII can be used in most situations. However, you can also use hexadecimal strings if:

- Your organization mandates their use.
- You have systems that do not support the use of ASCII strings.
- You want to use characters that are not supported in ASCII strings.

Examples of IPsec Configurations

The following examples are provided and depict several scenarios for using IPsec with your PS Series group. They provide configuration settings for the array and for initiators and hosts.

- Example 1: Transport mode (Host-to-Host) with certificates and PSK with Microsoft iSCSI Initiator
- Example 2: Tunnel Mode (between Linux hosts) using PSK
- Example 3: Tunnel Mode (between Linux hosts) using Certificate-Based Authentication

- Example 4: Tunnel Mode (Host-to-Gateway) using PSK with Cisco ASA Configuration

For information regarding connectivity considerations, limitations, and requirements for the various IPsec configurations, see IPsec Performance Considerations.

Example 1: Transport Mode (Host-to-Host) with PSK and IPv4

Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK illustrates a transport mode IPsec configuration in which one host is using IPv4 and PSK and another host is using IPv6 and certificates. Either IKEv1 or IKEv2 can be used in this configuration if supported by the host, but the example shown uses IKEv1.

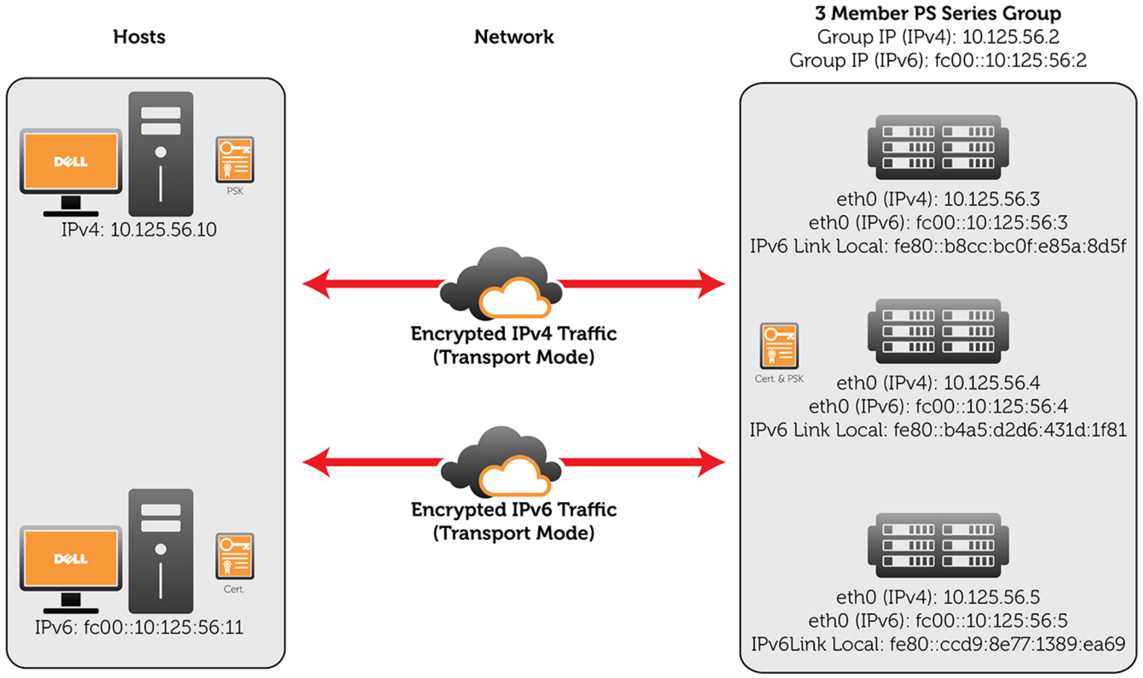
Hosts
Network
3 Member PS Series Group
Group IP (IPv4): 10.125.56.2
Group IP (IPv6): fc00::10:125:56:2
PSK
IPv4: 10.125.56.10
eth0 (IPv4): 10.125.56.3
eth0 (IPv6): fc00::10:125:56:3
IPv6 Link Local: fe80::b8cc:bc0f:e85a:8d5f
Encrypted IPv4 Traffic
(Transport Mode)
Cert. & PSK
eth0 (IPv4): 10.125.56.4
eth0 (IPv6): fc00::10:125:56:4
IPv6 Link Local: fe80::b4a5:d2d6:431d:1f81
Encrypted IPv6 Traffic
(Transport Mode)
Cert.
IPv6: fc00::10:125:56:11
eth0 (IPv4): 10.125.56.5
eth0 (IPv6): fc00::10:125:56:5
IPv6Link Local: fe80::ccd9:8e77:1389:ea69

Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK

iSCSI Initiator Configuration

In the example shown in Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK, the host systems' iSCSI initiators are configured using the Microsoft iSCSI Initiator. The following tables show how the initiator should be configured for IPv4 and IPv6 connections.

Table 18. iSCSI Initiator Configuration (IPv4) lists how the Microsoft iSCSI Initiator should be configured for the IPv4 connection shown in Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK.

Table 18. iSCSI Initiator Configuration (IPv4)

Setting	IPv4 Value
Rule Name	ToPSA_IPSEC_IPv4_CERT_IKEv1
Enabled?	Yes
Profiles	Domain,Private,Public
Mode	Transport
Endpoint1	10.125.56.10/32
Endpoint2	10.122.56.2/32 10.125.56.3/32

Setting	IPv4 Value
	10.125.56.4/32 10.125.56.5/32
Protocol	Any
Action	RequireInRequireOut
Auth1	ComputerCert
Auth1CAName	CN=sqaca
Auth1CertMapping	No
Auth1ExcludeCAName	No
Auth1CertType	Root
Auth1HealthCert	No
Anodes	DHGroup14-AES256-SHA384
QuickModeSecMethods	ESP:SHA1-AES256+60min+10000000kb

Table 19. iSCSI Initiator Configuration (IPv6) lists how the Microsoft iSCSI Initiator should be configured for the IPv6 connection shown in Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK.

Table 19. iSCSI Initiator Configuration (IPv6)

Setting	IPv6 Value
Rule Name	ToPSA_IPSEC_IPv6_PSK_IKEv1
Enabled?	Yes
Profiles	Domain,Private,Public
Type	Static
Mode	Transport
Endpoint1	fc00::10:125:56:11-fc00::10:125:56:11
Endpoint2	fc00::10:125:56:2-fc00::10:125:56:2 fc00::10:125:56:3-fc00::10:125:56:3 fc00::10:125:56:4-fc00::10:125:56:4 fc00::10:125:56:5-fc00::10:125:56:5 fe80::b8cc:bc0f:e85a:8d5f-fe80::b8cc:bc0f:e85a:8d5f fe80::b4a5:d2d6:431d:1f81-fe80::b4a5:d2d6:431d:1f81 fe80::ccd9:8e77:1389:ea69-fe80::ccd9:8e77:1389:ea69
Port1	Any
Port2	Any
Protocol	TCP
Action	RequireInRequireOut
Auth1	ComputerPSK
Auth1PSK	*password*

Setting	IPv6 Value
MainModeSecMethods	DHGroup14-AES256-SHA384
QuickModeSecMethods	ESP:SHA1-AES256+60min+10000000kb

CLI Commands (IPv4)

Enter the following CLI commands on the PS Series group to implement the IPv4 configuration shown in Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK:

```
> ipsec certificate load PSAcert IPsecPSA.pfx local password password
> ipsec certificate load RootCA rootca.cer root-ca
> ipsec security-params create RemPeer_CERT_Auth certificate id-type distinguished-name id-
value "CN=RemPeerDN"
> ipsec policy create ToRemPeer_IPv4_CERT_Ikev1 type v4 ip-addr 10.125.56.10 protocol any
action protect RemPeer_CERT_Aut
```

NOTE: The certificates must be loaded into the Local Computer CAPI store on the Windows host and uploaded to the / mgtdb/update folder on the PS Series array.

CLI Commands (IPv6)

Enter the following CLI commands on the PS Series group to implement the IPv6 configuration shown in Figure 10. Transport Mode (Host-to-Host) with Certificates or PSK:

```
> ipsec security-params create RemPeer_PSK_Auth pre-shared-key key password
> ipsec policy create ToRemPeer_IPv6_PSK_Ikev1 type v6 ip-addr fc00::10:125:56:11 protocol
tcp action protect RemPeer_PSK_Auth
```

NOTE: The certificates must be loaded into the Local Computer CAPI store on the Windows host and uploaded to the / mgtdb/update folder on the PS Series array.

Example 2: Tunnel Mode (Between Linux Hosts) Using PSK

In Figure 11. Tunnel Mode Between Linux Hosts Using PSK, an IPsec connection is established between Linux hosts running strongSwan and the PS Series group. The IPv4 and IPv6 traffic is protected using pre-shared keys (PSK). Either IKEv1 or IKEv2 can be used in this configuration. This particular example uses IKEv2.

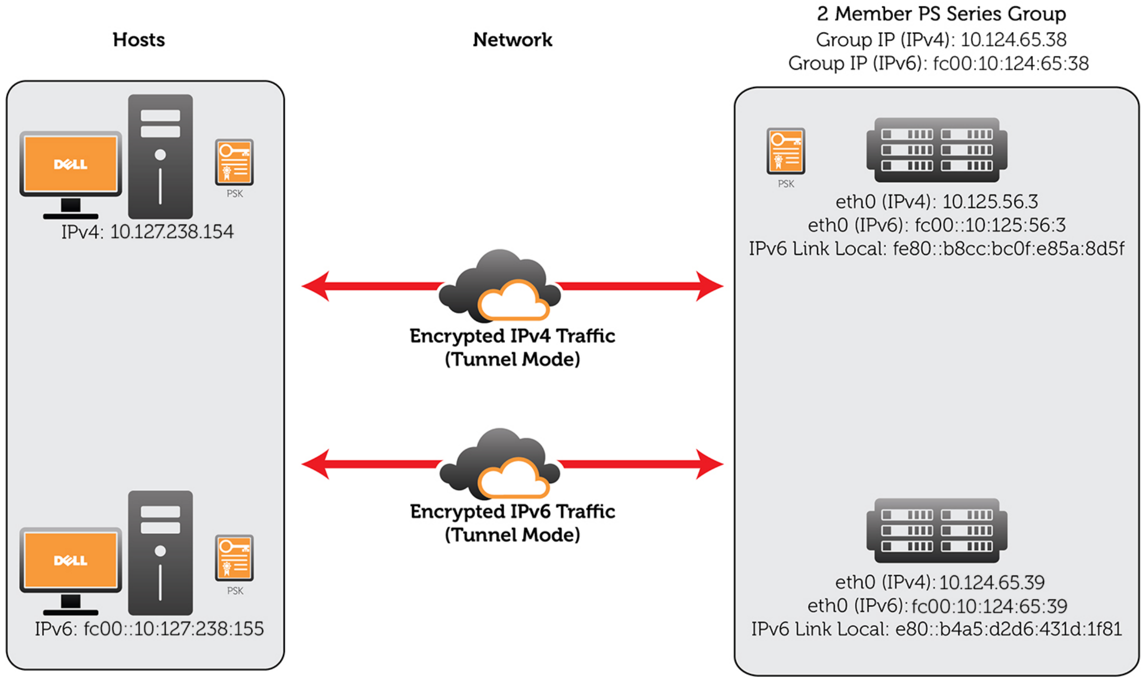
Hosts
Network
2 Member PS Series Group
Group IP (IPv4): 10.124.65.38
Group IP (IPv6): fc00:10:124:65:38
PSK
IPv4: 10.127.238.154
PSK
eth0 (IPv4): 10.125.56.3
eth0 (IPv6): fc00::10:125:56:3
IPv6 Link Local: fe80::b8cc:bc0f:e85a:8d5f
Encrypted IPv4 Traffic
(Tunnel Mode)
Encrypted IPv6 Traffic
(Tunnel Mode)
PSK
IPv6: fc00::10:127:238:155
eth0 (IPv4): 10.124.65.39
eth0 (IPv6): fc00:10:124:65:39
IPv6 Link Local: e80::b4a5:d2d6:431d:1f81

Figure 11. Tunnel Mode Between Linux Hosts Using PSK

iSCSI Initiator Configuration (IPv4)

This example uses the following configuration:

- Mint 17 (also known as Qiana)
- Linux Kernel 3.13.0-36-generic, x86_64
- strongSwan 5.1.2

The following configuration files are relevant:

- /etc/strongswan.conf is the configuration file that governs the operation of the strongSwan components (for example, debugging level, log file locations, and so on). You will not need to modify this file.
- /etc/ipsec.conf is the configuration file for IPsec that contains parametric information about the local host and the "peer" hosts that are configured to use IPsec.
- /etc/ipsec.secrets contains shared secrets (pre-shared keys).
- /etc/ipsec.d/cacerts/* contains Certificate Authority certificates.
- /etc/ipsec.d/certs/* contains intermediate and end-node certificates.

NOTE: Other directories within /etc/ipsec.d can hold certificate revocation lists and other kinds of special certificates, but they are not discussed here.

Initially, ipsec.conf contains exactly the following text:

```
# ipsec.conf - strongSwan IPsec configuration file

# basic configuration

config setup
        # strictcrlpolicy=yes
        # uniqueids = no

# Add connections here.

# Sample VPN connections

# conn sample-self-signed
#       leftsubnet=10.1.0.0/16
#      leftcert=selfCert.der
#      leftsendcert=never
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightcert=peerCert.der
#      auto=start

# conn sample-with-ca-cert
#      leftsubnet=10.1.0.0/16
#      leftcert=myCert.pem
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightid="C=CH, O=Linux strongSwan CN=peer name"
#      auto=start


Begin Pre-Shared Key Authentication, IPv4

1. strongSwan host IP address is 10.127.238.154
2. array addresses are 10.124.65.38 (the wka) and 10.124.65.39 (eth0)

The only two files modified on the strongSwan host are ipsec.conf and ipsec.secrets

# ipsec.conf - strongSwan IPsec configuration file

# basic configuration

config setup
        # strictcrlpolicy=yes
        # uniqueids = no

# Add connections here.

conn %default
    auto=route
    keyexchange=ikev1
    ike=3des-sha1-modp1024

conn kirt5eth0
    type=tunnel
    authby=psk
    right=10.124.65.39

conn kirt5wka
    type=tunnel
    authby=psk
    right=10.124.65.38
```

NOTE:

strongSwan allows you to specify properties that apply to all connections (`conn %default`). The `auto=route` directive tells strongSwan to install an IPsec security policy into the host's security policy database for every defined connection. If this directive were not present here, it would need to appear in the configuration for every connection.

`keyexchange=ikev1` is necessary because by default it will use/expect IKE version 1 for the key exchange algorithm. If you try to connect from the strongSwan side, strongSwan defaults to IKEv2 if this parameter is missing. Because IKEv1 is expected, the result is a failure to connect. The converse is not true; if the connection is initiated from the PS side, then strongSwan accepts either IKEv1 or IKEv2.

Two more "connections" are defined, one for each of the IP addresses on the array. The names are used by strongSwan to keep track of the connections. For example, you can ask for the status of a specific connection by name, shut down a connection by name, and so on. Connection names must be unique from one another. For example, you cannot have two connections named `kirt5eth0`.

`type=tunnel` tells strongSwan that tunnel-mode IPsec is to be used for the connection. The alternative is transport mode (`type=transport`). This mode must be consistent with the array's configuration; that is, you need to configure the array side to also use transport or tunnel mode.

`authby=psk` means that strongSwan expects to use pre-shared keys for authentication. In this example, it was implemented as a per-connection configuration item, but it could also be specified for the default connection, meaning that all connections would use pre-shared keys. The array's configuration must also use pre-shared keys.

Here is the `ipsec.secrets` file:

```
# This file holds shared secrets or RSA private keys for authentication.
# RSA private key for this host, authenticating it to any other host
# which knows the public part.  Suitable public keys, for ipsec.conf, DNS,
# or configuration of other implementations, can be extracted conveniently
# with "ipsec showhostkey".

: PSK "my_shared_key"
# 10.124.65.39 %any : PSK "my_shared_key"
# 10.124.65.38 %any : PSK "my_shared_key"
```

NOTE: This file defines a single pre-shared key that can be used for any connection defined in ipsec.conf. In the commented-out examples, connection-specific pre-shared keys are provided; `%any` represents any IP address on the strongSwan side. Using `%any` is easier than specifying that host's IP address (which if done incorrectly results in a failure to establish a connection).

Example 3: Tunnel Mode (Between Linux Hosts) Using Certificate-Based Authentication

In Figure 12. Tunnel Mode Between Linux Hosts Using Certificate-Based Authentication, an IPsec connection is established between Linux hosts running strongSwan and the PS Series group. The IPv4 and IPv6 traffic is protected using certificates. Either IKEv1 or IKEv2 can be used in this configuration. This particular example uses IKEv2.

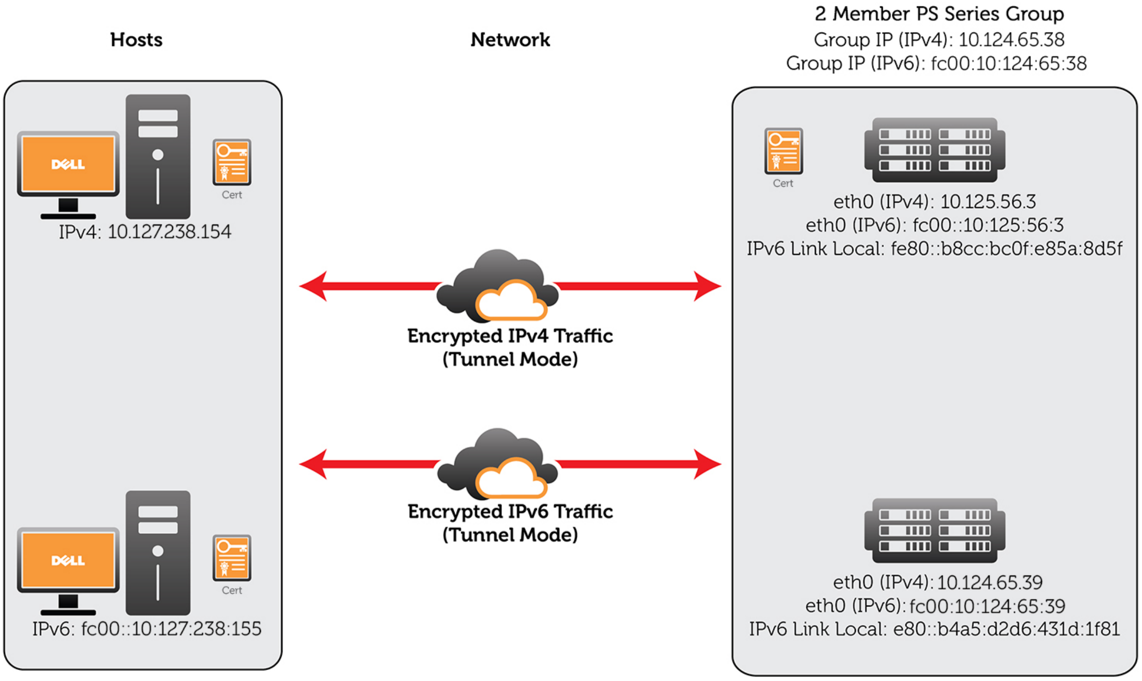
Hosts
Network
2 Member PS Series Group
Group IP (IPv4): 10.124.65.38
Group IP (IPv6): fc00:10:124:65:38
DELL
Cert
IPv4: 10.127.238.154
Cert
eth0 (IPv4): 10.125.56.3
eth0 (IPv6): fc00::10:125:56:3
IPv6 Link Local: fe80::b8cc:bc0f:e85a:8d5f
Encrypted IPv4 Traffic
(Tunnel Mode)
Encrypted IPv6 Traffic
(Tunnel Mode)
DELL
Cert
IPv6: fc00::10:127:238:155
eth0 (IPv4): 10.124.65.39
eth0 (IPv6): fc00:10:124:65:39
IPv6 Link Local: e80::b4a5:d2d6:431d:1f81

Figure 12. Tunnel Mode Between Linux Hosts Using Certificate-Based Authentication

iSCSI Initiator Configuration (IPv4)

This example uses the following configuration:

- Mint 17 (also known as Qiana)
- Linux Kernel 3.13.0-36-generic, x86_64
- strongSwan 5.1.2

The following configuration files are relevant:

- /etc/strongswan.conf is the configuration file that governs the operation of the strongSwan components (for example, debugging level, log file locations, and so on). You will not need to modify this file.
- /etc/ipsec.conf is the configuration file for IPsec that contains parametric information about the local host and the "peer" hosts that are configured to use IPsec.
- /etc/ipsec.secrets contains shared secrets (pre-shared keys).
- /etc/ipsec.d/cacerts/* contains Certificate Authority certificates.
- /etc/ipsec.d/certs/* contains intermediate and end-node certificates.

NOTE: Other directories within /etc/ipsec.d can hold certificate revocation lists and other kinds of special certificates, but they are not discussed here.

Initially, ipsec.conf contains exactly the following text:

```
# ipsec.conf - strongSwan IPsec configuration file

# basic configuration

config setup
      # strictcrlpolicy=yes
      # uniqueids = no

# Add connections here.

# Sample VPN connections

# conn sample-self-signed
#      leftsubnet=10.1.0.0/16
#      leftcert=selfCert.der
#      leftsendcert=never
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightcert=peerCert.der
#      auto=start

# conn sample-with-ca-cert
#      leftsubnet=10.1.0.0/16
#      leftcert=myCert.pem
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightid="C=CH, O=Linux strongSwan CN=peer name"
#      auto=start
```

Begin Certificate-Based Authentication, IPv4

1. strongSwan host IP address is 10.127.238.154
2. array addresses are 10.124.65.38 (the wka) and 10.124.65.39 (eth0)
3. 2048-bit RSA keys will be generated to encrypt/decrypt the local certificates (one for the array and one for the strongSwan host)
4. a self-signed root certificate will be generated
5. local certificate requests for both the array and the strongSwan client will be generated
6. certificate requests will be "signed" with our root certificate
7. the certificates and keys will be installed on the strongSwan host, then strongSwan will be reconfigured to use certificate-based authentication

Certificate Creation with OpenSSL:

1. Generate a 2048-bit RSA key. This is the "server" key, which will be used to generate a self-signed root certificate. Note that the minimum acceptable key length is 2048 bits:

```
1.1 draoidoir:fwoods> openssl genrsa -out server.key 2048

Generating RSA private key, 2048 bit long modulus
........+++.......................................................................+++
e is 65537 (0x10001)
```

2. With the server key in hand, generate a self-signed root certificate:

```
1.15 draoidoir:fwoods> openssl req -new -x509 -days 365 -key server.key -out root-ca.crt

You will be prompted to enter information that will be incorporated into the certificate
request. This is called a Distinguished Name or a DN. There are quite a few fields but you
can leave some blank
For some fields there will be a default value. If you enter '.', the field will be left
blank.
-----
Country Name (2 letter code) [AU]: US
State or Province Name (full name) [Some-State]: New Hampshire
Locality Name (eg, city) []: Nashua
Organization Name (eg, company) [Internet Widgits Pty Ltd]: Dell Equallogic
Organizational Unit Name (eg, section) []: Networking and iSCSI
Common Name (e.g. server FQDN or YOUR name) []: Joe Secure
Email Address []: Joe_Secure@dell.com
```

Now take a peek at the new root certificate:

```
draoidoir:fwoods> openssl x509 -text -noout -in root-ca.crt

Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 11801568908693661699 (0xa3c7986522fae803)
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=Joe Secure/emailAddress=Joe_Secure@dell.com
        Validity
            Not Before: Oct 14 19:01:25 2014 GMT
            Not After : Oct 14 19:01:25 2015 GMT
        Subject: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=Joe Secure/emailAddress=Joe_Secure@dell.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:c1:01:43:7f:63:96:ef:e2:c0:3c:9f:2c:0c:4e:
                    a6:73:e4:12:d3:cf:b1:2f:72:92:7b:67:1a:94:17:
                    08:09:83:ff:27:f6:fb:8b:95:4e:11:b1:3c:34:4b:
                    11:00:29:2b:95:50:e4:96:5a:ea:0a:73:8d:5d:09:
                    47:ac:88:b3:b3:6e:96:3a:29:ef:ff:6d:46:b6:3c:
                    5f:b6:68:05:af:f9:03:f5:52:30:fc:ce:94:30:3b:
                    08:98:1d:1c:9c:29:67:47:8d:2e:5c:c9:6f:1d:0a:
                    b7:92:1d:8a:9e:28:96:59:83:fa:5b:c6:0c:e1:75:
                    7d:09:d2:50:fa:c1:de:43:e8:62:df:fc:28:4f:6e:
                    21:ae:f8:2e:13:f6:e4:f0:6d:d6:49:0a:21:69:6c:
                    78:d3:25:dc:cf:d3:4f:bb:7e:c6:f1:74:17:b9:f7:
                    fd:66:e2:72:e5:02:0d:46:e2:9a:a2:06:70:a1:15:
                    af:dd:09:80:8b:e8:1b:c5:5f:0d:b6:17:47:cd:18:
                    8e:c5:69:a9:22:fd:2c:92:73:1f:e3:72:a2:8d:93:
                    fa:df:d8:a3:8f:fc:e0:69:35:62:80:cc:4b:07:34:
                    88:ac:c9:a9:0e:c0:34:a2:89:bb:11:27:b8:5e:64:
                    7a:2a:04:a5:59:76:a1:50:d3:c7:13:e8:de:63:e1:
                    75:fd
                Exponent: 65537 (0x10001)
            X509v3 extensions:
            X509v3 Subject Key Identifier:

7C:29:98:0E:A8:78:57:38:0C:88:4C:BA:6E:99:8B:55:64:B4:12:1D
            X509v3 Authority Key Identifier:
                keyid:7C:29:98:0E:A8:78:57:38:0C:88:4C:BA:6E:99:8B:55:64:B4:12:1D

            X509v3 Basic Constraints:
                CA:TRUE
    Signature Algorithm: sha256WithRSAEncryption
         2e:c1:56:89:b5:be:ad:d3:72:20:ba:76:6d:e3:35:3b:0e:3c:
         f7:5e:43:e8:b2:bc:e7:62:96:91:cd:64:ab:a5:39:74:8d:ab:
         a0:30:4d:78:af:19:73:19:fb:b6:18:0c:e9:70:fd:c4:30:26:
         c2:00:db:75:97:f6:19:0e:07:9e:01:96:e6:a9:a4:7f:0f:8f:
         2e:c3:96:b7:bb:b4:41:f7:48:d2:12:93:03:61:7a:91:53:49:
         97:24:80:f2:52:d0:ac:55:d3:f3:97:de:2a:22:26:db:7c:ff:
         1b:c9:b1:1f:4e:19:43:4b:99:9b:51:6e:f4:55:ed:89:9c:fe:
         d1:96:66:f8:5b:56:53:3d:dc:f5:ca:16:28:d7:5d:33:12:18:
         61:c8:9a:a2:16:0d:36:6c:a4:a2:ab:60:a8:ff:41:4f:63:de:
         83:2f:a1:b2:5e:5b:e6:c2:b7:23:5e:1e:d5:22:99:e1:50:93:
         3c:a8:29:db:cf:ff:f3:d2:65:ec:67:13:71:70:37:1e:e5:12:
         c9:5c:e0:76:b4:36:b5:b0:e0:59:42:81:d6:12:50:17:24:d6:
         34:06:82:22:08:0f:ea:fc:7b:83:e7:10:12:aa:00:10:f9:e3:
         cb:db:71:ab:99:45:3a:d0:07:b7:9c:12:e8:91:fa:63:d8:34:
         cd:70:24:47
```

3. Now generate a 2048-bit key for the two local certificates:

```
1.21 draoidoir:fwoods> openssl genrsa -out client.key 2048

Generating RSA private key, 2048 bit long modulus
...........................................+++.....+++
e is 65537 (0x10001)
```

4. Generate a "certificate request" for the array (kirt5):

```
1.31 draoidoir:fwoods> openssl req -key client.key -new -out kirt5.csr
```

```
You will be prompted to enter information that will be incorporated into the certificate
request. This is called a Distinguished Name or a DN. There are quite a few fields but you
can leave some blank
For some fields there will be a default value. If you enter '.', the field will be left
blank.
-----
Country Name (2 letter code) [AU]: US
State or Province Name (full name) [Some-State]: New Hampshire
Locality Name (eg, city) []: Nashua
Organization Name (eg, company) [Internet Widgits Pty Ltd]: Dell Equallogic
Organizational Unit Name (eg, section) []: Networking and iSCSI
Common Name (e.g. server FQDN or YOUR name) []: kirt5.lab.equallogic.com
Email Address []: Joe_Secure@dell.com

Please enter the following 'extra' attributes to be sent with your certificate request
A challenge password []:
An optional company name []:
```

5. Generate a certificate request for the strongSwan host:

```
1.32 draoidoir:fwoods> openssl req -key client.key -new -out draoidoir.csr

You are about to be asked to enter information that will be incorporated into your
certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value. If you enter '.', the field will be left
blank.
-----
Country Name (2 letter code) [AU]: US
State or Province Name (full name) [Some-State]: New Hampshire
Locality Name (eg, city) []: Nashua
Organization Name (eg, company) [Internet Widgits Pty Ltd]: Dell Equallogic
Organizational Unit Name (eg, section) []: Networking and iSCSI
Common Name (e.g. server FQDN or YOUR name) []: draoidoir.lab.equallogic.com
Email Address []: Joe_Secure@dell.com

Please enter the following 'extra' attributes to be sent with your certificate request
A challenge password []:
An optional company name []:
```

6. Sign the array's certificate request, creating the local certificate for kirt5:

```
1.33 draoidoir:fwoods> openssl x509 -req -in kirt5.csr -CA root-ca.crt -CAcreateserial -
CAkey server.key -out kirt5.crt -days 365

Signature ok
subject=/C=US/ST=New Hampshire/L=Nashua/O=Dell Equallogic/OU=Networking and iSCSI/
CN=kirt5.lab.equallogic.com/emailAddress=Joe_Secure@dell.com
Getting CA Private Key
```

7. Do the same for the strongSwan side:

```
1.34 draoidoir:fwoods> openssl x509 -req -in draoidoir.csr -CA root-ca.crt -CAcreateserial -
CAkey server.key -out draoidoir.crt -days 365

Signature ok
subject=/C=US/ST=New Hampshire/L=Nashua/O=Dell Equallogic/OU=Networking and iSCSI/
CN=draoidoir.lab.equallogic.com/emailAddress=Joe_Secure@dell.com
Getting CA Private Key
```

Now we can look at the two resulting certificates:

```
1.35 draoidoir:fwoods> openssl x509 -text -noout -in kirt5.crt
```

```
Certificate:
    Data:
        Version: 1 (0x0)
        Serial Number: 9335600219447230923 (0x818eb6effd3601cb)
    Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=Joe Secure/emailAddress=Joe_Secure@dell.com
        Validity
            Not Before: Oct 14 19:24:12 2014 GMT
            Not After : Oct 14 19:24:12 2015 GMT
        Subject: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=kirt5.lab.equallogic.com/emailAddress=Joe_Secure@dell.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:ef:67:f5:d5:06:06:38:33:54:41:44:7e:bc:6d:
                    70:35:ea:9a:10:7e:d4:f3:a2:c9:f5:3b:8c:35:19:
                    59:ba:77:09:01:b8:26:9e:e8:76:5e:54:06:82:5c:
                    f7:2c:a8:17:1a:16:bb:12:54:56:b5:3c:62:0b:58:
                    e8:4a:30:78:aa:3f:9f:9c:39:8a:3a:d2:9e:1d:3f:
                    dc:ea:4e:ff:e9:ae:a5:f0:c2:2c:ca:62:e2:56:00:
                    65:1b:96:0f:22:6a:c5:58:5c:00:d2:e3:b7:75:76:
                    02:1e:8e:47:59:07:8b:bc:4b:a5:b3:84:b0:ac:2e:
                    43:61:d2:29:a7:96:e2:60:21:5b:47:93:09:92:33:
                    7f:b9:94:78:6e:d3:cb:02:13:9d:18:53:62:f0:a2:
                    5a:27:c1:fd:31:8c:28:7a:48:8c:aa:5d:dc:6d:47:
                    dc:1b:90:60:f6:6d:67:6b:62:4a:05:23:9a:5b:72:
                    b0:fa:6d:d4:bd:40:0b:ab:6a:40:0e:85:c5:0d:90:
                    d4:a2:c0:9f:73:e8:13:a6:8b:9b:67:8f:15:5a:0c:
                    20:33:cc:90:7b:8a:7a:d9:af:18:03:c5:3e:bc:00:
                    a5:a4:71:ba:ab:d3:8d:85:17:44:e2:33:87:52:db:
                    b9:24:56:97:d6:62:40:13:82:a1:25:83:7c:1c:60:
                    52:93
                Exponent: 65537 (0x10001)
    Signature Algorithm: sha256WithRSAEncryption
         65:73:3a:88:4b:2a:8d:02:73:78:e0:84:f4:15:8e:7a:f3:85:
         8c:f6:76:3f:e9:17:c8:70:4d:24:93:64:8d:73:1b:35:4b:a5:
         39:32:0a:08:e4:a5:22:92:97:9c:96:ec:8d:f6:ae:42:f8:b2:
         e3:3d:f4:e2:af:47:a5:f3:58:65:59:93:66:18:29:3e:a0:8d:
         f6:6f:11:1e:c2:0f:48:2a:d8:ba:80:fe:08:48:b8:02:08:ee:
         16:5c:02:12:47:b3:b9:fe:56:ee:10:24:3d:cb:84:ff:70:cb:
         ea:52:bf:4f:08:c3:ef:2e:48:03:c2:2b:1c:68:32:01:ef:39:
         1d:58:ff:0f:00:ca:e1:2c:66:3f:fa:65:cf:ec:1d:25:ef:06:
         cc:7d:58:58:9f:0f:01:e6:98:c6:49:0a:6f:24:a4:03:7e:94:
         f9:cd:a5:6d:cb:8c:2b:25:ad:11:73:00:24:48:8c:a4:bd:73:
         72:eb:e5:d5:4d:b2:0b:43:51:24:7e:af:72:13:75:5a:08:8b:
         01:2f:c4:e5:5b:d9:67:5b:4f:a4:84:4d:94:1c:75:d8:19:76:
         a5:8c:e6:47:70:a1:da:22:78:6c:42:8f:d6:67:86:e5:bf:c2:
         a7:ba:7e:93:b1:75:b8:07:fc:08:4a:04:05:4f:22:70:49:db:
         65:b4:3a:13

1.36 draoidoir:fwoods> openssl x509 -text -noout -in draoidoir.crt

Certificate:
    Data:
        Version: 1 (0x0)
        Serial Number: 9335600219447230924 (0x818eb6effd3601cc)
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=Joe Secure/emailAddress=Joe_Secure@dell.com
        Validity
            Not Before: Oct 14 19:24:36 2014 GMT
            Not After : Oct 14 19:24:36 2015 GMT
        Subject: C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and
iSCSI, CN=draoidoir.lab.equallogic.com/emailAddress=Joe_Secure@dell.com
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:ef:67:f5:d5:06:06:38:33:54:41:44:7e:bc:6d:
                    70:35:ea:9a:10:7e:d4:f3:a2:c9:f5:3b:8c:35:19:
                    59:ba:77:09:01:b8:26:9e:e8:76:5e:54:06:82:5c:
                    f7:2c:a8:17:1a:16:bb:12:54:56:b5:3c:62:0b:58:
                    e8:4a:30:78:aa:3f:9f:9c:39:8a:3a:d2:9e:1d:3f:
                    dc:ea:4e:ff:e9:ae:a5:f0:c2:2c:ca:62:e2:56:00:
                    65:1b:96:0f:22:6a:c5:58:5c:00:d2:e3:b7:75:76:
                    02:1e:8e:47:59:07:8b:bc:4b:a5:b3:84:b0:ac:2e:
                    43:61:d2:29:a7:96:e2:60:21:5b:47:93:09:92:33:
                    7f:b9:94:78:6e:d3:cb:02:13:9d:18:53:62:f0:a2:
                    5a:27:c1:fd:31:8c:28:7a:48:8c:aa:5d:dc:6d:47:
                    dc:1b:90:60:f6:6d:67:6b:62:4a:05:23:9a:5b:72:
                    b0:fa:6d:d4:bd:40:0b:ab:6a:40:0e:85:c5:0d:90:
                    d4:a2:c0:9f:73:e8:13:a6:8b:9b:67:8f:15:5a:0c:
                    20:33:cc:90:7b:8a:7a:d9:af:18:03:c5:3e:bc:00:
                    a5:a4:71:ba:ab:d3:8d:85:17:44:e2:33:87:52:db:
                    b9:24:56:97:d6:62:40:13:82:a1:25:83:7c:1c:60:
                    52:93
                Exponent: 65537 (0x10001)
    Signature Algorithm: sha256WithRSAEncryption
         b6:c7:df:70:af:e3:19:af:6f:21:96:52:47:7f:4c:c3:c7:92:
         1e:bc:44:71:69:e7:7d:fc:17:27:91:6a:65:89:d0:09:13:6f:
         92:66:1e:5c:6a:48:66:32:ba:73:75:63:06:ee:a8:e9:e6:a4:
         f0:07:5e:84:1f:69:f1:cd:6f:4c:15:a4:9a:67:e0:85:94:41:
         5f:7b:fd:d9:e1:d9:42:56:88:5a:e1:70:77:ef:25:4b:df:0d:
         16:46:71:b9:62:2a:dc:47:a9:99:95:8a:31:c1:11:53:d4:7e:
         e0:2f:94:4a:c3:f1:96:53:b7:8e:9c:42:ae:3d:e9:4e:ca:a9:
         8a:10:48:5d:23:f2:92:0e:d2:c8:00:7c:5f:b0:ac:5f:54:76:
         74:e5:25:2c:1d:e3:57:93:4b:c0:89:b7:37:33:f9:86:1e:46:
         cf:6b:fc:b3:27:bd:66:6f:e9:52:6a:ac:c3:65:ba:25:a6:8d:
         45:72:ae:96:17:1f:91:8a:5f:a4:d8:6e:f0:be:62:da:21:57:
         05:d7:a0:41:3f:35:4b:75:7f:4a:3a:1c:21:80:b2:4d:db:19:
         da:8e:99:02:86:cf:e9:b1:f2:d6:2b:66:ed:e1:71:8d:3c:64:
         4e:b9:65:d5:72:49:2c:69:ce:1b:ed:30:54:88:2e:00:82:1d:
         7e:77:11:76
```

8. Push the root certificate and the array's certificate and private key up to the array:

```
1.38 draoidoir:fwoods> ftp 10.124.65.39

Connected to 10.124.65.39.
220 10.124.65.39 FTP server (NetBSD-ftpd 20100320) ready.
Name (10.124.65.39:fwoods): root
331 Password required for root.
Password:
230 User root logged in.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> prompt
Interactive mode off.
ftp> mput kirt5.crt root-ca.crt client.key
local: kirt5.crt remote: kirt5.crt
200 PORT command successful.
150 Opening BINARY mode data connection for 'kirt5.crt'.
226 Transfer complete.
1399 bytes sent in 0.00 secs (10045.7 kB/s)
local: root-ca.crt remote: root-ca.crt
200 PORT command successful.
150 Opening BINARY mode data connection for 'root-ca.crt'.
226 Transfer complete.
1497 bytes sent in 0.00 secs (10013.1 kB/s)
local: client.key remote: client.key
200 PORT command successful.
150 Opening BINARY mode data connection for 'client.key'.
226 Transfer complete.
1679 bytes sent in 0.00 secs (9759.8 kB/s)
ftp> bye
```

```
221-
    Data traffic for this session was 6250 bytes in 4 files.
    Total traffic for this session was 7728 bytes in 6 transfers.
221 Thank you for using the FTP service on 10.124.65.39.
```

9. Drop the certificates in place on the strongSwan host side:

```
# cp draoidoir.crt /etc/ipsec.d/certs
# cp root-ca.crt /etc/ipsec.d/cacerts
# cp client.key /etc/ipsec.d/private
```

10. Configure strongSwan to use the certificates for authentication. Here we have opted to use a Distinguished Name as the identifier on each side.
- On the strongSwan side, the identifier is the Distinguished Name that is contained in the certificate that will be presented by the array.
- On the array side we use the identifier that will be presented by the strongSwan host.
Each of these distinguished names are contained within the local certificates installed on each side:

```
# ipsec.conf - strongSwan IPsec configuration file

# basic configuration

config setup
        # strictcrlpolicy=yes
        # uniqueids = no

# Add connections here.

conn %default
    auto=route
    keyexchange=ikev1
    ike=3des-sha1-modp1024
    leftcert=draoidoir.crt
    leftsendcert=yes

conn kirt5eth0
    right=10.124.65.39
    type=transport
    #authby=psk
    authby=pubkey
    rightid="C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and iSCSI,
CN=kirt5.lab.equallogic. com, emailAddress=Joe_Secure@dell.com"

conn kirt5wka
    right=10.124.65.38
    type=transport
    #authby=psk
    authby=pubkey
    rightid="C=US, ST=New Hampshire, L=Nashua, O=Dell Equallogic, OU=Networking and iSCSI,
CN=kirt5.lab.equallogic. com, emailAddress=Joe_Secure@dell.com"

# Sample VPN connections

#conn sample-self-signed
#      leftsubnet=10.1.0.0/16
#      leftcert=selfCert.der
#      leftsendcert=never
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightcert=peerCert.der
#      auto=start

#conn sample-with-ca-cert
#      leftsubnet=10.1.0.0/16
#      leftcert=myCert.pem
#      right=192.168.0.2
#      rightsubnet=10.2.0.0/16
#      rightid="C=CH, O=Linux strongSwan CN=peer name"
#      auto=start

"leftcert=draoidoir.crt" tells strongSwan where it can find its local certificate (in /etc/
ipsec.d/certs).  This is the local certificate that it will present to the array.

"leftsendcert=yes" tells strongSwan that it should always send its certificate chain to any
peers.

"authby=pubkey" in each connection tells strongSwan that these peers will use certificate-
based authentication.

"rightid=..." is the identifier that strongSwan expects to see presented in the local
certificate that it receives from the array.


We also need to change the ipsec.secrets file:

# This file holds shared secrets or RSA private keys for authentication.

# RSA private key for this host, authenticating it to any other host
# which knows the public part.  Suitable public keys, for ipsec.conf, DNS,
# or configuration of other implementations, can be extracted conveniently
# with "ipsec showhostkey".

: RSA client.key
: PSK "my_shared_key"
# 10.127.238.154 10.124.65.39 : PSK "my_shared_key"
# 10.124.65.39 %any : PSK "my_shared_key"

": RSA client.key" tells strongSwan where to find the key file that is used
to decrypt the local certificate (in /etc/ipsec.d/private).
```

CLI Commands (IPv4)

Enter the following CLI commands on the PS Series group to implement the IPv4 configuration shown in Figure 12. Tunnel Mode Between Linux Hosts Using Certificate-Based Authentication:

```
> ipsec security-params create RemPeer_CERT_Auth_Tunnel certificate id-type distinguished-
name id-value
"CN=RemPeerDN" tunnel type v4 tun-ip-addr 10.125.56.10 require-ike-v2

>ipsec policy create ToRemPeer_IPv4_CERT_Ikev2 type v4 ip-addr 10.125.56.10 protocol any
action protect
RemPeer_CERT_Auth_Tunnel
```

NOTE: All certificates must be uploaded to the /mgtdb/update folder on the PS Series group.

CLI Commands (IPv6)

Enter the following CLI commands on the PS Series group to implement the IPv6 TCP configuration shown in Figure 12. Tunnel Mode Between Linux Hosts Using Certificate-Based Authentication:

```
> ipsec security-params create RemPeer_PSK_Auth_Tunnel pre-shared-key key <password> tunnel
type v6 tun-ip-addr
fc00::10:125:56:11 require-ike-v2 id-type domain-name id-value RemPeer.company.com

> ipsec policy create ToRemPeer_IPv6_PSK_Ikev2 type v6 ip-addr fc00::10:125:56:11 protocol
tcp action protect
RemPeer_PSK_Auth_Tunnel
```

Example 4: Tunnel Mode (Host-to-Gateway) Using PSK

In Figure 13. Tunnel Mode (Host-to-Gateway) Using PSK, a tunnel mode connection to a Cisco ASA gateway is established, using pre-shared keys to authenticate IPv4 traffic. The example uses IKEv1.

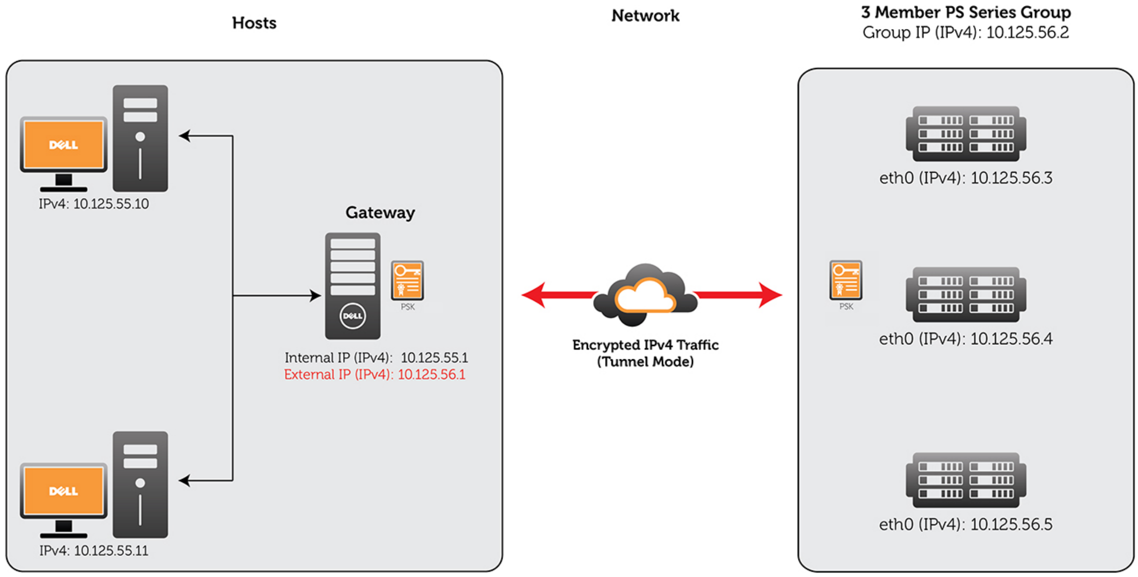

Figure 13. Tunnel Mode (Host-to-Gateway) Using PSK

.

Cisco ASA Configuration

The following Cisco ASA configuration for the gateway is shown in Figure 13. Tunnel Mode (Host-to-Gateway) Using PSK.

```
ASA Version 7.2(3)
!
hostname ciscoasa
domain-name company.com
enable password <> encrypted
names
!
interface Vlan10
nameif outside
security-level 0
ip address 10.125.56.1 255.255.255.0
!
interface Vlan200
nameif inside
security-level 100
ip address 10.125.55.1 255.255.255.0
!
interface Ethernet0/0
switchport access vlan 10
!
interface Ethernet0/1
!
interface Ethernet0/2
!interface Ethernet0/3
!
interface Ethernet0/4
!
interface Ethernet0/5switchport access vlan 200
!
interface Ethernet0/6
!
```

```
interface Ethernet0/7
!passwd <> encrypted
ftp mode passive
dns server-group DefaultDNS
domain-name company.com
access-list 101 extended permit ip 10.125.55.0 255.255.255.0 host 10.125.56.2
pager lines 24
mtu outside 1500
mtu inside 1500
icmp unreachable rate-limit 1 burst-size 1
asdm image disk0:/asdm-523.bin
no asdm history enable
arp timeout 14400
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout uauth 0:05:00 absolute
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
crypto ipsec transform-set aes_set esp-aes esp-sha-hmac
crypto map IPsecPSA 10 match address 101
crypto map IPsecPSA 10 set peer 10.125.56.2
crypto map IPsecPSA 10 set transform-set aes_set
crypto map IPsecPSA interface outside
crypto isakmp identity address
crypto isakmp enable outside
crypto isakmp policy 10
authentication pre-share
encryption aes
hash sha
group 2
lifetime 28800
crypto isakmp am-disable
telnet timeout 5
ssh timeout 5
console timeout 0
!
!
username name password <> encrypted

tunnel-group 10.125.56.2 type ipsec-l2l

tunnel-group 10.125.56.2 general-attributes

tunnel-group 10.125.56.2 ipsec-attributes

pre-shared-key *

no prompt

Cryptochecksum:<>

: end
```

This Cisco ASA configuration creates a secure connection to the group IP address. To establish secure connections to the individual network interfaces on each group member, you need to create an access list and crypto map for each interface. For example:

```
access-list <new ACL> extended permit ip 10.125.55.0 255.255.255.0 host <member physical
interface IP address>

crypto map IPsecPSAMem1Eth0 10 match address <new ACL>

crypto map IPsecPSAMem1Eth0 10 set peer <member physical interface IP address>

crypto map IPsecPSAMem1Eth0 10 set transform-set aes_set

crypto map IPsecPSAMem1Eth0 interface outside
```

CLI Commands (IPsec)

Enter the following CLI commands on the PS Series group to implement the configuration shown in Figure 13. Tunnel Mode (Host-to-Gateway) Using PSK:

```
> ipsec security-params create RemGW_PSK_Auth_Tunnel pre-shared-key key <password> tunnel
type v4 tun-ip-addr 10.125.56.1
> ipsec policy create ToRemGW_IPv4_PSK_Ikev1 type v4 ip-addr 10.125.56.0 netmask
255.255.255.0 protocol any action protect RemGW_PSK_Auth_Tunnel
```

IPsec Performance Considerations

The performance impact of IPsec varies by host and network configuration, and increases with the number of IPsec-protected iSCSI connections to the group. Even if IPsec is used only to protect traffic between group members, I/O performance is still affected. Based on these factors, you can expect that using IPsec might degrade I/O performance.

Although PS Series group members use hardware to accelerate cryptographic operations, many initiators perform these operations in software, which can cause a further reduction in the speed of communications between iSCSI initiators and the group.

IPsec Host Connectivity Considerations

- Enabling or disabling IPsec for the group using the ipsec enable and ipsec disable commands might disrupt host connectivity to the group for several minutes. To prevent unplanned outages, Dell recommends that IPsec be enabled or disabled during a planned maintenance window when volumes do not have any active iSCSI connections.
- Consult the documentation for your host operating systems, HBAs, and iSCSI initiators to verify that they support IPsec. The initiators' IPsec support might have known issues and idiosyncrasies that require additional planning or configuration.

  When configuring IPsec with Windows hosts, note the following limitations:

  - IPsec traffic is not always handled correctly if the IPsec policy is configured to protect only a subset of traffic between the host and the group. For example, if the IPsec policy protects only iSCSI traffic on port 3260, the Windows host might not perform reliably when connecting to the group. As a workaround, IPsec policies should apply to all traffic passing between the group and Windows systems. Microsoft KB article 2665206 discusses this workaround in greater detail.
  - IPsec must be configured using the Windows Firewall with Advanced Security. Do not use the IPsec option in the Microsoft iSCSI initiator, which does not have the capability to fully configure an IPsec configuration between the host and the group. Further, if you attempt to configure an IPsec connection using the iSCSI initiator, the system might not allow you to remove the partial configuration and replace it with a complete configuration created with Windows Firewall.
  - IPsec policies defined using the Local Security Policy Manager are not supported.

strongSWAN Limitations with IPsec

If you are using strongSWAN, the following limitations apply:

- If you are using certificates, the `uniqueids` keyword must be disabled (`uniqueids=no`).
- In rare cases, strongSWAN might negotiate standard frames in IPv6 environments even though jumbo frames are configured.
- If you are using IKEv2 and the certificate IDs are mismatched, the PSA might behave as if the security association (SA) has been established when it has not.
- strongSWAN does not create exceptions for IPv6 neighbor discovery in its Allow All IPsec policy. Consequently, neighbor discovery will fail and security associations (SA) will not be established. As a workaround, use an IPsec policy that uses ports and protocols to manage neighbor discovery.

IPsec Configuration Limitations

The following limitations apply when implementing IPsec:

- IPsec is supported only for certain PS Series array models, and can be enabled for a group only if all members support IPsec. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for more information.
- IPsec can be enabled and configured only with the Group Manager CLI. The Group Manager GUI provides no facility for configuring or monitoring IPsec.

- The PS Series firmware provides no mechanism for using IPsec to protect traffic between replication partners. It is technically possible to create IPsec polices on both the primary and secondary group in which each group treats the other as an iSCSI initiator and traffic is protected accordingly. However, this configuration is not supported, and Dell recommends against implementing it in a production environment.
- The PS Series array does not serve as an IPsec-secured gateway; it behaves as an IPsec-secured host only.
- You cannot use the save-config CLI command to preserve the group's IPsec certificates and pre-shared keys. The save-config command saves the CLI commands that were used to configure IPsec, but it does not save certificates that have been transferred to the array using FTP. Therefore, when you restore a configuration, you must manually restore any configuration options set using the ipsec certificate load, ipsec security-params create certificate, and ipsec security-params pre-shared-key commands.
- Kerberos-based authentication is not supported.
- Multiple Root Certificate Authorities (CA) are not supported.
- Certificate Revocation Lists (CRL) are not supported.
- Only users with group administrator privileges can configure IPsec.
- Perfect Forward Secrecy (PFS) is not supported.
- Encrypted private keys are not supported for X.509 format certificates.
- Dell recommends using a minimum of 3600 seconds and 10GB lifetime rekey values.
- IKE mobility is not supported.
- NAT Traversal (NAT-T) is not supported. Dell recommends against placing a firewall that performs address translation between the PS Series group and its IPsec peers.

Supported iSCSI Initiator Platforms

iSCSI initiators on the following hosts have been tested and verified for use with IPsec connections to PS Series groups:

- Microsoft Windows 2008, Windows 2008 R2, Windows 7, Windows Server 2012, and Windows Server 2012 R2
- Ubuntu Linux (using strongSWAN)

NOTE: Some Linux distributions use a different IKE implementation. For example: CentOS 6 uses Openswan. The configuration details change substantially depending on the IKE implementation used, and in particular, the examples provided in this document do not carry over to Openswan.

Requirements for IPsec Certificates

The following considerations apply to certificates:

- If a certificate that is uploaded to the array contains multiple Subject Alternative names, only the first name is used.
- Certificates can be imported using PKCS12 or X.509 formats.
- Encrypted private keys are not supported for X.509 format certificates. Use PKCS12 format certificates when encrypted private keys are required.
- The maximum supported certificate key size is 4096 bits, which applies to both local and root-CA certificates.
- Disabling support for legacy protocols prevents the following actions:
  - RSA-based SSH keys smaller than 2048 bits establishing SSH sessions to the group
  - All DSA-based SSH keys establishing SSH sessions to the group
  - Using the IKE (Diffie-Hellman) Key Exchange Group 2 algorithm
  - All IPSec certificates (both on the initiator and the group) using DSA keys establishing security associations
  - All IPSec certificates (both on the initiator and the group) with keys smaller than 2048 bits establishing security associations
  - Any certificate with keys smaller than 2048 bits from being imported into the group

Supported Relative Distinguished Names (RDN)

Table 20. Supported RDNs lists supported certificate Relative Distinguished Names (RDN).

Table 20. Supported RDNs

OID	RDN	Meaning
2.5.4.6	C	Country
2.5.4.7	L	Locality Name
2.5.4.5	serialNumber	Serial Number
2.5.4.9	STREET	Street Address
2.5.4.8	ST	State or Province
2.5.4.10	O	Organization
2.5.4.11	OU	Organizational Unit
2.5.4.3	CN	Common Name
	DC	Domain Component
RSA.1.9.1	MAILTO	PKCS9 Email Address
RSA.1.9.1	emailAddress	PKCS9 Email Address
RSA.1.9.2	unstructuredName	PKCS9 Unstructured Name
2.5.4.4	SN	Surname
2.5.4.12	title	Title
2.5.4.41	Name	Name
2.5.4.42	givenName	Given Name
2.5.4.43	initials	Initials
2.5.4.44	generationQualifier	Generation Qualifier
2.5.4.45	x500UniqueIdentifier	X.500 Unique Identifier
2.5.4.46	dnQualifier	DN Qualifier
RSA.1.9.8	unstructuredAddress	PKCS9 unaddr
2.5.4.65	pseudonym	Pseudonym

IPsec Supported Authentication Algorithms

Table 21. Supported IPsec Authentication Algorithms lists supported IPsec authentication algorithms.

Table 21. Supported IPsec Authentication Algorithms

Algorithm Type	Supported Algorithms
IKE Integrity	• HMAC-SHA1–96 • HMAC-SHA2–224 • HMAC-SHA2-256 • HMAC-SHA2-384 • HMAC-SHA2-512
IKE Encryption	• 3DES-CBC • AES-CBC • AES-CBC–192

Algorithm Type	Supported Algorithms
	- AES-CBC–256
IKE (Diffie-Hellman) Key Exchange	- 2 (if legacy support is not disabled) - 14 - 24
IPsec Integrity	- HMAC-SHA1–96 - HMAC-SHA2–224 - HMAC-SHA2-256 - HMAC-SHA2-384 - HMAC-SHA2-512
IPsec Encryption	- NULL - 3DES-CBC - AES-CBC - AES-CBC–192 - AES-CBC–256

NOTE: IKE (Diffie-Hellman) Key Exchange Group 2 algorithm is supported only if legacy support is not disabled.

Requirements for Pre-Shared Keys

Pre-shared keys that are used with the group must meet the following requirements:

- A text string of 6 to 64 printable ASCII characters, meeting these specifications:
  - Letters and numbers are allowed, but letters with accent marks, such as é, ç, ñ, or ü, are not.
  - Spaces are not allowed.
  - ASCII keys cannot begin with “0x” or “0X”; that prefix is reserved for hexadecimal keys.
  - You can use the following nonalphanumeric characters: ! " # $ percent & ` ( ) * + , - / : ; < = > ? @ [ \ ] ^ _ ` { | } ~ .
- An even number of hexadecimal digits, meeting these specifications:
  - The string must be 12 to 128 ASCII characters long.
  - The string must be preceded by either `0x` or `0X`. The prefix does not count toward the number of characters and is not part of the pre-shared key. The system will interpret any pre-shared key that does not begin with `0x` or `0X` as a text string, even if it contains only hexadecimal characters.

Protect Network Traffic with IPsec

To enable IPsec protection for traffic between the group and iSCSI initiators, use the following basic process:

NOTE: This process is not required for protecting communications between group members. After IPsec is enabled, all network traffic between group members is automatically protected, without need for further configuration.

1. A group administrator creates security parameters to specify how traffic should be authenticated.
2. A group administrator creates policies to identify traffic and determine what action to take for it:
   - Traffic is dropped.
   - Traffic is allowed to pass directly through to the array in the clear.
   - Traffic is protected using certificates or pre-shared keys.

NOTE: IPsec configurations cannot be modified. They must be removed and then recreated using the new configuration.

Protect Communication Between Group Members

To enable IPsec security for communication between group members, use the **ipsec enable** CLI command. No further configuration actions are required.

Protect iSCSI Initiator Connections

IP traffic between the group and iSCSI initiators is not automatically protected after IPsec has been enabled. Configure an IPsec configuration as follows:

NOTE: See the *Dell EqualLogic Group Manager CLI Reference Guide* for command syntax and examples of the CLI commands.

1. If you are authenticating with certificates rather than pre-shared keys, load local and root-CA certificates using the **ipsec certificate load** command. (See About IPsec for more information.)
2. Create a security parameter using one of the **ipsec security-params create** commands, based on the authentication method:
   - If you are using a certificate, use **ipsec security-params create certificate**.
   - If you are using a pre-shared key, use **ipsec security-params create pre-shared-key**.

   NOTE: Local and root-CA certificates must be loaded before you can create certificate-based security parameters. This step is not required for security parameters using pre-shared keys.
3. Create a policy that defines a particular set of network traffic and applies a specific action to that traffic, a process that is conceptually similar to creating a firewall rule. You can either drop the traffic, allow it to pass through, or protect it using a security parameter.

   NOTE: If you are creating policies that drop traffic or allow it to pass, you do not have to create the corresponding security parameter.
4. Perform additional host or initiator configuration tasks required to use IPsec. See your operating system or iSCSI initiator documentation for instructions.
5. IPsec must be enabled for the IPsec configuration to take effect. However, you can still create IPsec configurations while IPsec is disabled.

Add Members to an IPsec-Enabled Group

You can add new group members to an existing IPsec-enabled group, provided the new member is a model that supports IPsec. See the documentation for the **setup** command in the *Dell EqualLogic Group Manager CLI Reference Guide* for instructions on joining an IPsec-enabled group.

Some older PS Series array models do not support IPsec. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for more information.

Remove an IPsec Configuration

IPsec configurations cannot be modified. They must be removed and then recreated using the new configuration.

NOTE: This process might disrupt connections that are covered by the policy being deleted. Dell recommends changing the configuration during a maintenance window.

If you are removing an IPsec configuration, you must delete the components in the reverse order in which they were applied:

1. Delete the policy.
2. Delete the security parameter.
3. Delete the certificate (if certificate-based protection is being used).

This same ordering rule applies when deleting any component used in an IPsec configuration. To delete a security parameter, you must first delete any policies using it. To delete a certificate, you must first delete any security parameters that use it.

About Dedicated Management Networks

For increased security, or if your environment requires the separation of management traffic and iSCSI traffic, you can configure a dedicated management network (DMN) that is used only for administrative access to the group. The management network is separate from the network that handles iSCSI traffic to the group.

- Without a dedicated management network (the default configuration), administrators connect to the group IP address for both administrative access to the group and iSCSI initiator access to iSCSI targets (volumes and snapshots).
- With a dedicated management network, administrators do not use the group IP address for administrative access to the group. Instead, administrators connect to the management network address. All iSCSI traffic – including traffic by replication partners, and access to Dell EqualLogic Auto-Snapshot Manager/Linux Edition (ASM/LE), Dell EqualLogic Auto-Snapshot Manager/Microsoft Edition (ASM/ME), and Dell Virtual Storage Manager for VMware (formerly ASM/VE) – continues to use the group IP address. SAN Headquarters can connect to the group using either the management network address or the iSCSI address.

NOTE: Dedicated management networks are supported with both IPv4 and IPv6. You can configure IPv4 dedicated management networks using the Group Manager GUI; however, IPv6 networks must be configured with the CLI.

All currently available PS Series array control modules have an additional network interface dedicated for use with a management network. However, some older array models do not have this extra network interface. On these systems, if you enable the dedicated management network:

- The highest-numbered interface becomes the DMN interface and is no longer available for iSCSI traffic.
- Your iSCSI bandwidth might decrease because iSCSI bandwidth traffic is limited to the remaining interfaces.

See the hardware documentation for your array model for more information.

If the management interface fails in a single-member group, or if the management network loses connectivity, you lose management access to the group. However, you can always connect to the serial port on a group member and use the Group Manager CLI to manage the group.

Using IPv6 with Management Networks

When using IPv6 addressing in your management network, note the following limitations.

NOTE: Dedicated management networks are supported with both IPv4 and IPv6. You can configure IPv4 dedicated management networks using the Group Manager GUI; however, IPv6 networks must be configured with the CLI.

- Management network are supported through IPv6 addressing only when all members of a group are using PS Series firmware version 9.0 or later.
- If you use a management port, configure it promptly after adding the new member. That way, the management address remains accessible when the group needs to recover from failures.
- FS Series appliances do not support IPv6.
- Virtual volumes (VVols) do not support IPv6.
- Management network support for IPv6 must be configured through the CLI.
- If IPv4 and IPv6 addresses are configured for MWKA (management well-known addresses), you can use the following methods to display the IP addresses:
  - To display IPv4 addresses, click **Group** → **Group Configuration** → **General**. The addresses are shown in the General Settings panel. (IPv6 addresses are blank.)
  - To display IPv6 addresses, use the CLI **grpparams management-network show** command. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

Configure a Management Network

Perform these tasks before configuring a management network:

- Make sure your network environment can support a dedicated management network. You need a subnet for the management network that is separate from the subnet (or subnets) for iSCSI traffic.

- Obtain an IP address and default gateway information for the management network address. This address is the one to which administrators can connect.
- For each group member, obtain an IP address for the management network interface. The IP address must be on the same subnet as the management network address, and this subnet should not be the same as the one used for data I/O.
- On each group member, connect at least one network interface, other than the highest-numbered interface, to the iSCSI network and configure and enable the interface. For the best performance, connect, configure, and enable all iSCSI-eligible network interfaces. To support control module failover, connect the ports on the active and secondary control modules to the network.
- On each group member, connect the highest-numbered network interface on the active and secondary control modules to the management network.

  For example, if you have a three-port control module, connect the Ethernet 2 port on both control modules to the management network. For some control modules, this interface is labeled `Management`.
- If any non-array devices on the SAN network have Link Layer Discovery Protocol (LLDP) enabled, the switch must have LLDP support enabled. If the SAN switch does not provide LLDP support, disable LLDP on all non-array devices connected to the switch. For instructions on disabling LLDP on your devices, refer to the user manual of the device.

To configure a management network:

1. Click **Group → Group Configuration**.
2. Click the **Advanced** tab.
3. Click **Configure management network** to open the Configure Dedicated Management Network dialog box.
4. Select **Enable dedicated management network**.
5. Type the management network IP address in the **Management IP address** field.
6. Type the default gateway in the **Default gateway** field.
7. For each group member:
   a. Double-click the member name to configure and enable at least one network interface, other than the highest-numbered interface, on the iSCSI network. For best performance, connect, configure, and enable all iSCSI-eligible interfaces.
   b. Click **Configure for management-only access** next to the highest-numbered network interface.
   c. In the **Modify IP settings – management network** dialog box:
      - Type an IP address that is on the management network subnet.
      - Type a subnet mask (netmask).
      - Select **Enable this interface**.
      - If it is available (does not appear dimmed), select **Restrict to management access**.
      - Click **Yes** in the confirmation message box.

After setting up the management network on each group member, verify the network configuration in the Configure Management Network dialog box.

When you configure a management network correctly, the highest-numbered interface is on the same subnet as the management IP address and `Mgmt` appears in the Traffic column. The remaining interfaces for iSCSI traffic are on a different network and `SAN` appears in the Traffic column.

1. Click **OK** to complete the dedicated management network configuration.
2. Click **Yes** in the Warning dialog box to restart the Group Manager GUI session using the new management IP address.

When you complete the management network configuration, administrators cannot log in to the group using the group IP address. Instead, administrators must use the new management IP address. Any open GUI or CLI sessions using the group IP address eventually time out and close.

After configuring a dedicated management network, you might need to:

- Inform administrators of the new management network IP address.
- If you run the Group manager GUI as a standalone application and have a shortcut on the computer's desktop, the group address in the shortcut is not updated with the new management address. You must uninstall and then reinstall the GUI application.

- If you are running SAN Headquarters, you must update the group IP address in the application to the dedicated management address. For more information, see the SAN Headquarters documentation.
- If you are using an NTP server, Dell recommends that the NTP server be on the same subnet as the dedicated management network.

Display Management Network Information

1. Click **Group** → **Group Configuration**.
2. Click the **Advanced** tab.
3. Click **Configure management network** to display details.

The management network address is shown in the **Management IP address** field. (This field is not shown if the system does not have a management network.)

The General Settings panel also shows the group IP address, which you use for all iSCSI traffic, including traffic between replication partners.

Unconfigure a Management Network

You can unconfigure a dedicated management network and reenable the group IP address to be used for group management.

This operation is in two parts. First, you disable and stop using the current management network address, and then you change the management network interfaces on all the group members so that they again support all traffic.

Disable and Stop Using the Management Network Address

1. Click **Group** → **Group Configuration**.
2. Click the **Advanced** tab.
3. Click **Configure management network**.
4. Clear **Enable dedicated management network** in the Configure Dedicated Management Network dialog box.
5. Click **OK**. The Group Manager GUI automatically restarts, using the group IP address.
6. Log in to the group.

Reconfigure the Member Management Network Interfaces

For each member in the group, prepare the physical connections:

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. Click the **Network** tab.
4. Select the former management interface and then click **Modify IP settings**.
5. If you do not want to use the interface, delete the IP address and click **OK** in the Modify IP Settings dialog box.

   If you want to use the interface:

   a. Change the IP address and subnet mask to the iSCSI network.
   b. Clear **Restrict to management access**.

      NOTE: On some arrays, you cannot select Restrict to management access. In such cases, this option is not bold.

   c. Select **Enable this interface**, then click **OK**.

To log in to and manage the group, connect to the group IP address.

Enable Secure Erase

You can securely erase data so that it cannot be recovered.
When a secured volume is deleted, the following occurs:

- The volume is not preserved in the recycle bin.

- Pages are overwritten by a random pattern as a background operation.
- Page map entries are cleared.

To enable Secure Erase:

1. Click **Group** → **Group Configuration**.
2. Click the **Advanced** tab.
3. In the Secure Erase panel, check the **Enable secure erase data** checkbox.

7

About Volume-Level Security

To secure your data, you must prevent access by unauthorized iSCSI initiators. By controlling access to your iSCSI targets, you can secure access to individual volumes. Group Manager provides several ways to control access to your volumes. You can use these security measures in tandem with group-level and NAS-level security to provide the required level of security for your data.

- You can specify a CHAP user name, IP address, or iSCSI initiator name. This information is used by the access method that applies to a volume and its snapshots. You can then use a CHAP account on an external RADIUS authentication server to authenticate iSCSI targets in a PS Series group.
- You can allow or disallow initiators with different iSCSI qualified names (IQN) access to a volume and its snapshots.
- You can use an iSNS (Internet Storage Name Service) server for initiator discovery of iSCSI targets.
- You can set permissions for a volume as either read-write (default) or read-only.

Connect Initiators to iSCSI Targets

To access iSCSI targets (volumes and snapshots) in a PS Series group, you must install an industry-standard iSCSI initiator on a computer. An example of an industry-standard iSCSI initiator is the one that is built in to Microsoft Windows.

NOTE: Access to iSCSI targets is through TCP port 3260 (the standard iSCSI port).

See your initiator documentation for the exact procedure for logging in to an iSCSI target.

In general, to log in:

1. Specify the group IP address as the discovery address or target portal in the iSCSI initiator configuration interface. If you are using iSNS, the initiator automatically discovers targets from the iSNS server that you configured in the group.

   The initiator displays a list of iSCSI targets from the group.
2. Log in to a target. The initiator must match at least one of the target’s access control policies.

   As part of the login procedure, you might need to enter a CHAP user name and password (secret) and target authentication credentials.

After the initiator logs in to the iSCSI target, the computer sees the target as a disk that you can format using the usual operating system utilities. You can then partition the disk and create a file system as needed.

NOTE:

- **In some file systems, volumes and snapshots must have read-write permission even if the file system is read-only.**
- **Both hardware and software iSCSI initiators are available from a variety of vendors. Install and configure an initiator using the vendor-supplied instructions. See the *Dell EqualLogic PS Series Storage Arrays iSCSI Initiator and Operating System Considerations* document for more information about iSCSI initiator configuration or contact Dell Technical Support.**

Access Control Methods

Access control methods determine which hosts and clusters can connect to which volumes while simultaneously preventing unauthorized access to iSCSI target volumes and snapshots. Access methods restrict access to iSCSI target volumes and snapshots to specified initiators, restricted by CHAP user name, iSCSI initiator name, or IP address. The access method can contain one or more of these restrictions.

Different access methods are available depending on the needs of your environment:

- An access policy consists of a set of extended access points. Each extended access point enables users to provide a set of access attributes describing the endpoints, such as an IQN initiator name, CHAP name, and IP addresses. After an access policy is associated with a volume, all the endpoints described by the extended access points will have access to the volume.
- An access policy group is a set of access policies that can be associated to a volume. When an access policy group is associated with a volume, all endpoints described within those access policies have access to the volume.
- A basic access point provides the traditional direct method for connecting a single endpoint to a single volume. Basic access points cannot be reused, transferred, or shared with other volumes. They are associated directly with the volume to which they provide access, and if that volume is deleted, the basic access point is also deleted.

All of these access methods can be used with each other to fulfill the particular needs of your computing environment.

The main distinction between access policies and basic access points is that access policies exist independent of the volumes to which they provide access. This flexibility offers the following significant advantages over traditional basic access points:

- Access policy reuse — A single access policy can now be associated with multiple volumes and all the access attributes are specified only once, requiring less manual configuration and less possibility of data entry errors.
- Single point of change — Changes to access attributes, IP address, initiator IQN name, or CHAP user name are all specified at one place and not repeated across volumes. You can now change these attributes at one place, and the changes are instantly propagated to all volumes using those access policies.

About Access Policies

In earlier versions of the PS Series firmware, security protection was accomplished by individually configuring an access control record for each volume to which you wanted to secure access. Each volume supported up to 16 different access control records, which together constituted an access control list (ACL). However, this approach did not work well when large numbers of volumes were present. To address that issue, Group Manager incorporates access policies and access policy groups that can be applied to one or more volumes.

Each access policy lets you specify one or more of the following authentication methods:

- CHAP user name (Challenge Handshake Authentication Protocol)
- IP address
- iSCSI initiator name

When you create a volume, you can assign it to an existing access policy, which determines which hosts will have access to that volume. In addition, you can allow or disallow volume access from multiple initiators, depending on your configuration needs.

An access policy can apply to the volume, its snapshots, or both. For example, you can authorize computer access to a volume and its snapshots or only to the volume.

Access Policies: Use Cases

The following use cases show different ways of working with access policies.

Study 1: Grant volume access to a single host using its iSCSI Initiator Name or IP address

Scenario: A group administrator wants to grant volume access to a host that is using a software iSCSI initiator. The admin wants to define access using either the iSCSI initiator name or IP addresses.

Solution:

1. Click **Group** → **Group Configuration**.
2. Click the **Access Policies** tab. In the Access Policies panel, create an access policy for the host.
3. Add an access point to the access policy that specifies either the initiator name or a list of IP addresses for each of the interfaces on the host.
4. With the access policy selected, go to the Targets panel and click **Add**.
5. Select the name of the volume to which you want to grant access and click **OK**.

Study 2: Apply an existing policy to a volume

Scenario: A user wants to grant access to a volume using a previously specified access policy (or policy group), without having to reenter the IP address, initiator name, and CHAP user name.

Solution:

If the volume has not been created yet:

1. Run the **Create volume** wizard to define the parameters of the new volume. Complete wizard steps 1 and 2.
2. When the Define iSCSI Access Points step is reached, select **Select or define access control policies**.
3. Specify the access policies that you want to associate with the new volume and finish the wizard as usual. The new volume will be created with the specified access policies enabled.

If the volume already exists:

1. Click **Group** → **Group Configuration**.
2. Click the **Access Policies** tab.
3. In the Access Policies panel, select the access policy or policy group that you want to use.
4. Go to the Targets panel and click **Add**.
5. Select the name of the volume to which you want to grant access and click **OK**.

Study 3: Change the CHAP/Initiator/IP Address

Scenario: The group administrator has granted a host access to a set of volumes by associating an access policy for the host to a set of volumes. Now the admin wants to change the CHAP user name (or the initiator name or the IP address) without reconfiguring the access policies for each volume.

Solution:

1. Create an access policy for host A that is populated with all of the host's details.
2. Associate that policy with the volumes to which access is granted.
3. To change the host details, modify this single access policy with the new information. The changes are automatically carried across each volume to which the policy is associated.

Study 4: Grant a cluster shared volume access to all nodes in the cluster

Scenario: The group administrator wants to grant access so that all nodes in the cluster can see the set of cluster-shared volumes.

Solution:

Create an access policy for each node in the cluster to describe the node's access attributes.

1. Build an access policy group that includes all of the access policies.
2. With the access policy group selected, go to the Targets panel and click **Add**.
3. Select the name of the cluster volumes to which you want to grant access and click **OK**.

Study 5: Add or remove a new node from a group assigned to multiple volumes

Scenario: The group administrator has configured a set of volumes to be accessible from cluster A. If a new node is added or removed from the cluster, the group administrator should be able to configure this node without adding or removing access policies for each of the volumes.

Solution:

1. Create an access policy group for cluster A that associates it to the volumes that the cluster nodes can access.
2. To add a new node to cluster A, define a new access policy for the node and associate it with the group policy. This association instantly makes the new node part of the group.

3. To remove a node from cluster A, remove the node's access policy from the group policy. This disassociation instantly removes the node from the group.

Study 6: Determine each volume a host/cluster can access

Scenario: The group administrator needs to determine which volumes a host can access.

Solution:

1. Click **Group → Group Configuration**.
2. Click the **Access Policies** tab.
3. In the **Access Policies** panel, select either the policy group or the access policy that is assigned to the host or cluster.
4. Refer to the **Targets** panel to view the volumes that are associated with the selected policy.

Create a New Access Policy

Access policies associate one or more endpoints to any number of available volumes. The access policy is created independently of any specific volume association and then associated with the selected volumes as a separate operation.

1. Click **Group → Group Configuration**.
2. In the Group Configuration panel, click the **Access Policies** tab to view the Access Policies panel.
3. In the Access policies subpanel, click **New**.
4. In the New Access Policy dialog box, specify a name for the new access policy and (optionally) a description.

   NOTE: Policy names can be up to 31 characters long. Space and underscores are not permitted. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. In the New Access Policy window, click **New**.
6. In the New Extended Access Point dialog box, define an access point by specifying your host's authorization information. This information can be a CHAP account name, an iSCSI initiator name, an IP address, or any combination of these parameters.

   - To specify a new IP address, click **Add** next to the IP Addresses list. When the text cursor appears in the box, type the IP address.
   - To modify an existing IP address, double-click it or select the line and click **Modify**.
   - To delete an IP address from the list, select it from the list and click **Delete**.

   NOTE:

   - **You can specify up to 16 unique IP addresses within a single access point.**
   - **Asterisk characters can be used within an IP address to indicate that any value is accepted in an octet (for example: 12.16.*.*).**
7. Click **OK** to save the access point, which now appears in the Extended Access Points list.

   - To add another access point to this policy, click **Add** and repeat the previous step. You can associate up to 16 access points with an access policy.
   - To modify the currently selected access point, click **Modify**.
   - To remove the currently selected access point, click **Delete**.
8. Click **OK** to save the new policy, which now appears in the Access Policies list.

Create a New Basic Access Point

Basic access points connect a specific endpoint to a specific volume. Unlike access control policies and access policy groups, basic access points are always created in reference to a specific volume; the basic access point cannot be shared or reassigned to other volumes. When a volume is deleted, all the basic access points associated with it are also deleted.

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name on which you want to assign the access point.

NOTE: Click Access to show all access policies and access points that are currently assigned to this volume.

3. In the Activities panel, click **Add basic access point**.
4. Specify the authorization parameters (CHAP account name, iSCSI initiator name, or IP address) for the host that you are configuring.

   NOTE: Asterisk characters can be used within an IP address to indicate that any value is accepted in an octet (for example: 12.16.*.*). Group Manager displays *.*.*.* if you leave the IP address field empty to indicate that any IP address will be accepted.

5. In the **Access point applies to** section, specify whether the access point should apply to volumes and snapshots, volumes only, or snapshots only.
6. Click **OK** to save the new access point to the volume.

Modify or Delete a Basic Access Point

The following considerations apply:

- You cannot change the iSCSI initiator name in a basic access point. Instead, you must delete the access point and recreate it using the new initiator name.
- You must also verify and, if necessary, change the access control information for any computer that you want to authorize for access to targets. Computers that met the original access conditions might be unable to log in to the target.

To modify a basic access point:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click the **Access** tab to open the Access Control List window.
4. Select the access point and click **Modify**.
5. Change the information as needed and click **OK**.

To delete a basic access point:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click the **Access** tab to open the Access Control List window.
4. Select the access point and click **Remove**.
5. Confirm that you want to remove the access point.

Modify Access Policies and Basic Access Points by Volume

You can make modifications to any policy group, access policy, or access point on a per-volume basis, including:

- Adding additional groups, policies, or points to the volume
- Editing an existing access policy or access point
- Disassociating an access group, access policy, or access point from the volume

NOTE: Be aware that any changes that are made to an access policy or access policy group are propagated to all volumes that are associated with that policy.

Disassociating an access policy or an access policy group does not delete that policy; it removes its association with that volume. The policy remains in effect on any other volumes to which it is associated, and remains available for assignment to other volumes. However, when a basic access point is removed from a volume, it is permanently deleted.

To modify a volume's access policies or basic access points:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name on which you want to modify access attributes.
3. Click the **Access** tab to open the Access Control List panel. This panel displays the access policy groups, access policies, and basic access points that are assigned to this volume.

4. Perform the desired action in the corresponding subpanel:

   Add, Modify, or Remove an Access Policy Group:

   - To bind an additional access policy group to the volume, click **Add** to open the Add Access Policy dialog box. You can then select additional groups that you want to bind to this volume.
   - To make changes to the access policies within an access policy group, select a group policy and click **Modify** to open the Edit Access Policy Group dialog box. You can add, edit, or remove the access policies assigned to this group.
   - To disassociate an entire policy group from the volume, select that policy group name and click **Delete**. When prompted to confirm the decision, click **Yes**.

   Add, Modify, or Remove an Access Policy:

   - To bind an additional access policy to the volume, click **Add** to open the Add Access Policy dialog box. You can then select the additional access policies that you want to bind to this volume.
   - To make changes to the access points within an access policy, select a policy and click **Modify** to open the Edit Access Policy dialog box. You can create new access points, edit existing access points, or remove access points that belong to this policy.
   - To disassociate an access policy from the volume, select the policy name and click **Delete**. When prompted to confirm the decision, click **Yes**.

   Add, Modify, or Remove a Basic Access Point

   - To associate an additional access point to the volume, click **New** to open the New Basic Access Point dialog box. You can then define an additional access point.
   - To change the parameters of an existing access point (CHAP name, iSCSI name, or IP address), select the access point that you want to edit and click **Modify**.
   - To disassociate a basic access point from the volume, select the access point name and click Delete. When prompted to confirm the decision, click **Yes**.

Associate Access Control Policies with Volumes

Associating access control policies (including access control groups) to volumes can be achieved in two ways, depending on how many access policies you want to assign to how many volumes. You can select a single policy and choose which volumes you want to assign it to, or you can select a single volume and choose which policies you want to associate with it. If you are associating only one policy to one volume, either method works.

Associate a Single Access Policy with One or More Volumes

1. Click **Group** → **Group Configuration**.
2. Click the **Access Policies** tab.
3. In the **Access Policies** panel, select the access policy group or the access policy that you want to bind to a volume.
4. In the Targets panel, click **Add**. The Pick Volumes dialog box opens and displays a list of the volumes to which the policy has not already been associated.
5. Select the checkbox next to each volume that you want to associate with the selected access policy.
6. In the **Applies to** section, specify whether the access policy should apply to volumes and snapshots, volumes only, or snapshots only.
7. Click **OK**. The volumes that you selected now appear in the Targets panel, indicating that they are now associated with the currently selected access policy.

Associate One or More Access Policies with a Single Volume

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name to which you want to provide access.
3. Click the **Access** tab.
4. In the Access Control List panel, click the **Add** button that is associated with the access policies.

   The Add Access Policies dialog box displays a list of access policies that are not yet associated with the selected volume.
5. Select the checkbox next to each policy (up to 4) that you want to associate with the selected volume and click **OK**.

Create an Access Policy Group

Access policy groups combine individual access policies together so that they can be managed as a single entity.

1. Click **Group → Group Configuration**.
2. Click the **Access Policies** tab.
3. In the Access Policies panel, locate the Access Policy Groups section and click **New** to open the New Access Policy Group dialog box.
4. Specify a policy name for the new group and (optionally) a description.
5. In the Access Policies section, click **Add**. The Add Access Policy dialog box opens and displays a list of all the access policies that have been previously defined.
6. Select the checkbox next to the name of each access policy that you want to include in the group. Up to 64 policies can be selected.

   NOTE: You have the option to create additional access policies from this dialog box, as well as modify or remove any of the existing access policies.

   - To create a new access policy, click **New**.
   - To edit an existing access policy, select the checkbox next to that policy name and click **Modify**.
   - To delete an existing access policy from the group, select the checkbox next to that policy name and click **Delete**.
7. Click **OK**. The Create Access Policy Group dialog box now shows the names of the access policies that you selected to be part of the new policy group.

   NOTE: You can modify or remove these access policies.
8. Click **OK** to finish and create the new access policy group.

Associate an Access Policy Group to a Volume

Associating access policy groups to volumes can be achieved in two ways, depending on how many access policy groups you want to assign to how many volumes. You can either select a single access policy group and choose which volumes you want to associate it with, or you can select a single volume and choose which policy groups you want to associate it with. If you need to associate only one access policy group to one volume, either method works.

Associate a Single Access Policy Group to One or More Volumes

NOTE: The access policy group must already exist before performing this procedure.

1. Click **Volumes**.
2. In the Activities panel, click **Manage access policies**. The Access Policies panel opens.
3. In the **Access policy groups** subpanel, select the access policy group that you want to associate with a volume.
4. In the Targets panel, click **Add**. A list of available volumes (not already assigned to the group) is displayed in the Pick Volumes dialog box.
5. From the list of volumes, select the checkbox next to each volume that you want to associate with the group policy.
6. In the **Applies to** section, specify whether the access policy should apply to volumes and snapshots, volumes only, or snapshots only.
7. Click **OK**. The volumes that you selected now appear in the Volumes Using Policy subpanel, indicating that they are now associated to the selected access policy group.

Associate One or More Access Policy Groups to a Single Volume

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name to which you want to provide group access.
3. Click the **Access** tab.
4. In the Access Control List panel, click **Add** to add the access policy group.

   The Add Access Policy Groups dialog box displays a list of access policy groups that are not yet associated with the selected volume.

5. Select the checkbox next to each group policy name that you want to associate with the selected volume and click **OK**.

Manage Access Controls for VDS/VSS Access

To allow VDS and VSS access to the group, you must create at least one VDS/VSS access control policy that matches the access control credentials you configure on the computer by using Remote Setup Wizard or Auto-Snapshot Manager/Microsoft Edition. The same access control constructs (access policies, access policy groups, and basic access points) are available for defining VDS/VSS access.

1. Click **Group → Group Configuration**.
2. Click the **VDS/VSS** tab.
3. Take the appropriate action in the **VDS/VSS Access Control List** panel to either add, modify, or remove access.

   Add, Modify, or Remove an Access Policy Group:

   - To add an access policy group for VDS/VSS access, click **Add**. Select the checkbox next to the additional policy groups that you want to assign and click **OK**.
   - To make changes to the access policies within an access policy group, select a group policy and click **Modify** to open the Edit Access Policy Group dialog box. You can add, modify, or remove the access policies within this group.
   - To remove an entire policy group from VDS/VSS access, select that policy group name and click **Delete**. When prompted to confirm the decision, click **Yes**.

   Add, Modify, or Remove an Access Policy:

   - To add an additional access policy for VDS/VSS access, click **Add**. Select the checkbox next to the additional access policies that you want to assign and click **OK**.
   - To make changes to the access points within an access policy, select a policy and click **Modify** to open the Edit Access Policy dialog box. You can create new access points, edit existing access points, or remove access points that belong to this policy.
   - To remove an access policy from VDS/VSS access, select the policy name and click **Delete**. When prompted to confirm the decision, click **Yes**.

   Add, Modify, or Remove a Basic Access Point:

   - To create an additional access point for VDS/VSS access, click **New** to open the New Basic Access Point dialog box. You can then define an additional access point.
   - To change the parameters of an existing access point (CHAP name, iSCSI name, or IP address), select the access point that you want to edit and click **Modify**.
   - To remove a basic access point from VDS/VSS access, select the access point name and click **Delete**. When prompted to confirm the decision, click **Yes**.

Authenticate Initiators with CHAP

CHAP (Challenge Handshake Authentication Protocol) is a network login protocol that uses a challenge-response mechanism. You can use CHAP to authenticate iSCSI initiators by specifying a CHAP user name in an access control policy. To meet this condition, a computer must supply the user name and its password (or “secret”) in the iSCSI initiator configuration interface when logging in to the target.

Using CHAP for iSCSI authentication can help you manage access controls more efficiently because it restricts target access by using user names and passwords, instead of unique IP addresses or iSCSI initiator names.

Before you can use CHAP for initiator authentication, you must set up the CHAP accounts consisting of a user name and password (or “secret”). Two options are available for accounts; you can use both options simultaneously in a group:

- CHAP accounts in the group

  Local CHAP accounts do not rely on any external system. You can create up to 100 local CHAP accounts.
- CHAP accounts on an external RADIUS authentication server

Using a RADIUS server to manage CHAP accounts is helpful if you are managing a large number of accounts. However, computer access to targets depends on the availability of the RADIUS server.

NOTE: If you use CHAP for initiator authentication, you can also use target authentication for mutual authentication, which provides additional security.

Display Local CHAP Accounts

To display local CHAP accounts:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.

   The Local CHAP Accounts panel lists all current CHAP accounts.

   NOTE: Starting with firmware v9.1.x, the CHAP password is no longer displayed in clear text format.

Create a Local CHAP Account

CHAP accounts are a method of ensuring that only authorized users can access a PS Series group. You can create local CHAP accounts or you can use a RADIUS server.

Before you create an account:

- You can decide whether to verify iSCSI initiator credentials against local CHAP accounts first (before verifying external CHAP accounts on a RADIUS server).
- You need the following information:
  - CHAP user name
  - Password (otherwise known as a CHAP *secret*). For optimal security, passwords must contain at least 12 characters (preferably random). Individual iSCSI initiators have their own rules and restrictions for length and format. Consult your initiator documentation for details.

To create a local CHAP account:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. (Optional) Select **Enable local authentication and check local first** in the iSCSI Authentication panel.
4. In the Local CHAP Accounts panel, click **Add** to open the Add CHAP Account dialog box.
5. Type a CHAP user name and, optionally, a password.
   - The user name can be up to 63 printable characters (any characters except space and colon).

     NOTE: If the user name contains a pound-sign character, enclose the name in quotation marks (for example, "chap#user"). Otherwise, the system will read the characters after the pound sign as a comment and not include them in the user name.
   - The password can be up to 255 printable characters (any characters except space and colon). If you do not enter a password, the group automatically generates a password that is 16 characters long).
6. Select whether to enable the account. You must enable an account to use it for initiator authentication. You can modify an account and enable or disable it later.
7. Click **OK**.
8. Click **Save all changes**.

   NOTE: In the iSCSI initiator authentication area, you can select Enable RADIUS authentication for iSCSI initiators, Consult locally defined CHAP accounts first, or both. Make sure that *at least* one of these choices is selected. If neither option is selected, the PS Series group will lock out all iSCSI initiator logins.

After creating the CHAP account, you can:

- Create an access control policy and use the CHAP user name in the policy

- Enable target authentication (for mutual authentication)

Modify a Local CHAP Account

To modify a local CHAP account:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the Local CHAP Account panel, select the account name that you want to edit and click **Modify**. The Modify CHAP Account dialog box opens.
4. Change the name or password or enable or disable the account, as needed.
5. Click **OK**.

Delete a Local CHAP Account

To delete a local CHAP account:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the Local CHAP Account panel, select the CHAP account that you want to remove and click **Delete**.
4. Confirm the deletion by clicking the **Yes** button.

Configure CHAP for Initiator Authentication on Existing Volumes

To configure CHAP for an existing volume:

1. Click **Volumes**
2. Expand **Volumes** and then select the volume that you want to configure.
3. Click the **Access** tab.
4. In the Basic access points section, click **New** to open the New Basic Access Point dialog box.
5. In the dialog box, type a description for the volume and a CHAP account name.

   Names can be up to 63 ASCII characters.
6. Select whether the access point applies to volumes and snapshots, volumes only, or snapshots only.
7. Click **OK**.

Configure CHAP for Initiator Authentication on New Volumes

To configure CHAP for a new volume:

1. Click **Volumes**.
2. In the Activities panel, click **Create volume** to open the Create Volume dialog box.
3. In the dialog box, type the general and space information for the volume.
4. For the iSCSI access information, select **Define one or more basic access points**.
5. Click **Add** to open the New Basic Access Point dialog box.
6. In the dialog box, type a description for the volume and a CHAP account name.

   Names can be up to 63 ASCII characters.
7. Select whether the access point applies to volumes and snapshots, volumes only, or snapshots only.
8. Click **OK** and finish typing the information for the volume.

Configure CHAP Accounts on a RADIUS Authentication Server

To use a CHAP account on an external RADIUS authentication server for iSCSI initiator authentication:

1. Set up the RADIUS server and CHAP accounts. (The RADIUS server must be accessible to all the group members.)

2. Click **Group** → **Group Configuration**.
3. Click the **iSCSI** tab.
4. In the iSCSI Authentication panel, select **Enable RADIUS authentication for iSCSI initiators**.
5. (Optional) Select **Enable local authentication and check local first**.
6. Click **RADIUS settings** to configure the group to use a RADIUS server (if you have not already done so).
7. Add at least one RADIUS server by clicking the **RADIUS settings** button and adding the IP address of the RADIUS authentication server.
8. Click **OK** to save the changes.

After creating the CHAP account, create an access control policy for a volume and specify the CHAP user name in the policy.

NOTE: In the iSCSI Authentication panel, you can select either Enable RADIUS authentication for iSCSI initiators, Enable local authentication and check local first, or both. Make sure that *at least* one of these options is selected. If neither option is selected, the PS Series group will lock out all CHAP logins.

Configure Target Authentication

If you configure initiator authentication though a local CHAP account or a CHAP account on a RADIUS authentication server, you can also allow the iSCSI initiator to authenticate iSCSI targets in a PS Series group. The combination of initiator and target authentication is called mutual authentication and provides additional security.

With target authentication, when the initiator tries to connect to a target, the target supplies a user name and password to the initiator. The initiator compares the user name and password to mutual authentication credentials that you configure in the initiator configuration interface. The iSCSI connection succeeds only if the information matches.

A group automatically enables target authentication using a default user name and password, which you can change. Whether the initiator requires target authentication depends on the initiator configuration settings.

To display the current target authentication user name and password:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the iSCSI Authentication panel, click **Modify**. The Modify Target CHAP Account dialog box opens.
4. In the dialog box, type the target authentication user name and password.
5. Click **OK** to save the changes.

About iSNS Servers

In a shared storage environment, you must control computer access to iSCSI targets (volumes and snapshots), because multiple computers writing to a target in an uncoordinated manner might result in volume corruption.

When an initiator tries to log in to a target, the group uses access control policies to determine if access should be authorized. However, access control policies do not prevent multiple initiators, either on the same computer or different computers, from accessing the same target.

Therefore, by default, the group disables multihost (shared) access to a target. Only one iSCSI qualified name (IQN) can connect to a target at one time.

If all group members are not running PS Series firmware version 5.0 or later, the group allows multihost access to targets.

An iSNS (Internet Storage Name Service) server can facilitate iSCSI initiator discovery of iSCSI targets in a SAN.

Prerequisites for Configuring an iSNS Server

The following considerations might apply when configuring an iSNS server:

- By default, the group disables target discovery by iSNS servers. If you want iSNS servers to discover a target, you must enable this functionality on the target. Set up the iSNS server and configure the iSCSI initiator to use the iSNS server for discovery. See your iSNS server and iSCSI initiator documentation for details.
- The iSNS server must be accessible to all the group members.
- A group disables automatic discovery of group targets by iSNS servers only if all the group members are running PS Series firmware version 5.0 or later. If a member is running a previous firmware version, iSNS servers can automatically discover all group targets.
- You can specify up to three IP addresses. The group uses only one iSNS server at a time. The first server listed is the default server. If the default server is not available, the group uses the other servers in the order in which they were specified. Click the Up and Down arrows to change the IP address order.
- If you are using a port number other than the default of 3205, you need an IP address for each iSNS server in the format `nnn.nnn.nnn.nnn:port`.

Enable or Disable iSNS Discovery

You can configure a PS Series group to use an iSNS (Internet Storage Name Service) server, which facilitates the discovery of iSCSI targets in the group.

By default, the group disables automatic discovery of targets by iSNS servers. If you want an iSNS server to automatically discover a group target, you must enable this functionality on the target.

Requirement
A group disables automatic discovery of group targets by iSNS servers only if all the group members are running PS Series firmware version 5.0 or later. If a member is running a previous firmware version (for example, if you downgrade a member from v5.0), iSNS servers can automatically discover all group targets.

You cannot use the Group Manager GUI to enable or disable iSNS discovery for a volume or snapshot. Instead, you must use the following CLI command lines:

```
volume select volume_name isns-discovery enable | disable
volume select volume_name snapshot select snapshot_name isns-discovery enable | disable
```

Configure an iSNS Server in the Group

To configure an iSNS server:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the iSCSI Discovery panel, click **Add** to open the iSNS Server dialog box.
4. Specify the IP address and optional port for an iSNS server and click **OK**. (The default port setting is 3205.)
5. Repeat steps 3 and 4 to add up to two additional servers.

NOTE: To prevent hosts from discovering iSCSI targets that do not have the correct CHAP credentials, select the Prevent unauthorized hosts from discovering targets checkbox.

Modify an iSNS Server

To modify the IP address for an iSNS server:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the iSCSI Discovery panel, select the address and click **Modify**.
4. Change the IP address as needed.

5. Click **OK**.
6. Click **Save all changes**.

Delete an iSNS server

To delete the IP address for an iSNS server to remove the server from the configuration:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the **iSCSI Discovery** panel, select the server's IP address.
4. Click **Delete**.
5. When prompted to confirm the decision, click **Yes**.

Prevent Discovery of Unauthorized Targets

By default, iSCSI initiators that use discovery try to log in to group targets protected by CHAP, even if they do not have the correct access credentials. These attempts can result in a large number of events logged in the group and is not an efficient use of resources.

You can prevent computers from discovering unauthorized targets by enabling the iSCSI discovery filter. If you enable the iSCSI discovery filter, initiators discover only those targets for which they have the correct access credentials.

iSCSI Access Requirements

To access an iSCSI target (for example, a volume or snapshot), an iSCSI initiator must meet the security requirements identified in Table 22. Access Requirements for iSCSI Targets.

Table 22. Access Requirements for iSCSI Targets

Security Condition	Description
Network access	To discover targets, the initiator must have network access to the group IP address.
Initiator access controls	(Optional) If the initiator enabled target authentication (sometimes called mutual authentication), the target authentication credentials in the group must match the credentials that were configured in the initiator. These credentials apply to all group targets.
Target access controls	The initiator must meet all the conditions in one access control policy for the target. (See About Multihost Access to Targets.)

Enable the iSCSI Discovery Filter

To enable the iSCSI discovery filter:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the iSCSI Discovery panel, select **Prevent unauthorized hosts from discovering targets**.

Disable the iSCSI Discovery Filter

To disable the iSCSI discovery filter:

1. Click **Group** → **Group Configuration**.
2. Click the **iSCSI** tab.
3. In the iSCSI Discovery panel, clear **Prevent unauthorized hosts from discovering targets**.

About Multihost Access to Targets

In a shared storage environment, you must control computer access to iSCSI targets (volumes and snapshots), because multiple computers writing to a target in an uncoordinated manner will result in volume corruption.

When an initiator tries to log in to a target, the group uses access control policies to determine if access should be authorized. However, access control policies do not prevent multiple initiators, either on the same computer or different computers, from accessing the same target. By default, the group disables multihost (shared) access to a target. Therefore, only one iSCSI qualified name (IQN) can connect to a target at one time.

Restriction: If all group members are not running PS Series firmware version 5.0 or later, the group allows multihost access to targets.

If you disable multihost access to a volume, when an initiator tries to log in to the volume:

- If an iSCSI initiator is not connected to the volume, the group uses access control policies to determine whether to authorize access.
- If an initiator is connected to the volume, the group compares the IQN of the current connection to the IQN of the incoming connection. If the IQNs are not the same, access is denied. If the IQNs are the same, the group uses access control policies to determine whether to authorize access.

However, some environments might need multihost access to a target. You can enable multihost access to a target if you meet one of the following conditions:

- Your cluster environment gives the initiators on each cluster computer a different IQN, and the environment can manage multiple connections to a target. For example, the environment uses a Distributed Lock Manager or SCSI reservations.
- Your multipathing solution does not use the same IQN on all initiators, and you cannot modify the names to be the same.
- You use an environment, such as a virtual server, that can manage multiple connections to the same iSCSI target (for example, through SCSI reservations).
- Initiators on a single computer do not use the same IQN.

In all cases, use access control policies as the primary method of protecting iSCSI targets in a group.

You can enable or disable multihost access when creating a volume. You can also modify a volume or snapshot and enable or disable multihost access.

Allow or Disallow Multihost Volume Access

In a shared storage environment, you must control computer access to iSCSI targets (volumes and snapshots), because multiple computers writing to a target in an uncoordinated manner will result in volume corruption.

You can allow or disallow multihost (shared) access to a volume. If you disallow multihost access to a volume, only one iSCSI qualified name (IQN) can connect to the volume at one time. However, if you have a certain environment, you might want to allow multihost access to a volume.

NOTE: Before disallowing multihost access to a volume, disconnect all initiators from the volume except one, unless the initiators have the same IQN. If multiple initiators with different IQNs have connections to the volume, you cannot disallow multihost access.

To allow or disallow multihost access to a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Set access type** to open the Set Access Type dialog box.
4. Choose to allow or disallow multihost access.
5. Click **OK**.

About Snapshot Access Controls

Online snapshots are seen on the network as iSCSI targets. It is important to protect your snapshots from unauthorized and uncoordinated access by iSCSI initiators.

NOTE: When a snapshot is online and accessible, a user or application can change the contents of the snapshot. If the content changes, the snapshot no longer represents a point-in-time copy of a volume and has limited use for data recovery.

All iSCSI target security mechanisms apply to snapshots, including access control policies, which prevent unauthorized iSCSI initiator access to a volume and its snapshots.

About Multihost Snapshot Access

In a shared storage environment, you must control computer access to iSCSI targets (volumes and snapshots), because multiple computers writing to a target in an uncoordinated manner can result in volume corruption.

You can allow or disallow multihost (shared) access to a snapshot. If you disallow multihost access to a snapshot, only one iSCSI qualified name (IQN) can connect to the snapshot at one time. However, if you have a certain environment, you might want to allow multihost access to a snapshot. See About Multihost Access to Targets.

NOTE: To disable multihost access to a snapshot, first disconnect all initiators from the snapshot except one. If multiple initiators have connections when you try to disable multihost access, the operation fails unless the initiators have the same IQN.

Allow or Disallow Multihost Snapshot Access

To enable or disable multihost access to a snapshot:

1. Click **Volumes**.
2. Expand **Volumes** and then expand the volume name.
3. Select the snapshot timestamp.
4. In the Activities panel, click **Set access type** to open the Set Access Type dialog box.
5. To allow multiple initiators to access the target, select the **Allow simultaneous connections from initiators with different IQNs** checkbox. (By default, this checkbox is not selected.)
6. Click **OK**.

About NAS Container Security

You control access to your NAS containers through volume-level and group-level security.

Windows and UNIX operating systems use different mechanisms for user identification, authentication, and resource access control. The file security style controls the type of operations that are permitted in the NAS container.

When you create a NAS container, the NAS cluster applies the NAS clusterwide default file security style. When a file or directory is created, the default NAS container security style, which controls the permissions and ownership, is applied.

You can modify a NAS container to change the file security style. The modification will affect only those files and directories that are created after the modification.

A NAS cluster supports the following security styles:

- UNIX — Controls file access using UNIX permissions in all protocols. A client can change a permission only by using the chmod and chown commands on the NFS mount point. You can specify UNIX permissions for files and directories created in the NAS container by Windows clients.

Windows clients cannot change any file or directory permissions. Read, write, and execute access is controlled by the UNIX permissions for Windows files and directories, which you set in Group Manager.

- NTFS — Controls file access by Windows permissions in all protocols. A client can change the permission and ownership by using the Windows Security tab. This security style is the default style.

  All access permissions are controlled by the Windows administrator through the use of access control lists or share-level permissions.
- Mixed — Supports both NTFS and UNIX security styles. The permissions and ownership for a file or directory will be the last ones set. Permissions and access rights are automatically translated from one protocol to another.

  A Windows user can override UNIX user settings, and a UNIX user can override Windows user settings.

In multiple protocol environments, it can be beneficial to set UNIX security style for UNIX clients and set NTFS security style for Windows clients.

Modify the File Security Style for a NAS Container

To modify the file security style for a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Modify Settings**.
3. In the Modify Settings dialog box, click the **Access Permissions** tab.
4. Select the file security style: NTFS, UNIX, or mixed. The default security style is NTFS.

   If you select UNIX for the security style, the Modify Settings dialog box displays additional fields that enable you to specify the UNIX file and directory permissions for the NAS container.
5. Click **OK**.

Modify the UNIX Permissions for Windows Directories or Files

If you are using the UNIX file security style, you can specify the UNIX permissions for directories that are created in the NAS container by Windows clients using SMB.

To modify the UNIX directory permissions for a NAS container:

1. Click **NAS**, expand **NAS cluster** and **Local Containers**, and then select the NAS container name.
2. Click **Modify Settings**.
3. In the Modify Settings dialog box, click the **Access Permissions** tab.
4. Set the permission (Read, Write, or Execute) for each user (Owner, Group, and Others).

   Alternatively, you can specify the equivalent three-digit format in the **Numeric value** field.
5. Click **OK**.

PS Series Group Operations

You can perform basic and advanced operations on the PS Series SAN.

Table 23. PS Series Group Operations

	Group Operations	Pool Operations	Volume Operations
Basic	Add a member Modify the IP address or name Set or convert a RAID policy Convert a RAID policy	Create, modify, or delete a pool	Create, modify, or delete volumes Create, modify, or delete volume collections Create, modify, or delete volume folders
Advanced	Shut down a group Override automatic load balancing	Merge storage pools	Restore deleted volumes Change the reported volume size Reclaim unallocated space Set a volume or snapshot with lost blocks online Manage storage capacity utilization on demand (thin provisioning) Improve pool space utilization (template volumes and thin clones) Improve performance with data center bridging (DCB)

About Group Network Configuration

The group network configuration, which you set when creating a group, includes the group name and group IP address.

You can modify the group name and IP address, although you should change them only when necessary.

- The group name identifies the group on the network. Group names must be unique in your network environment. A group name can consist of 1 to 54 ASCII characters (letters, numbers, or hyphens). The first ASCII character must be a letter or number.
- The group IP address is the network address for the group. You use the group IP address as the iSCSI discovery address when connecting initiators to iSCSI targets in the group.

  You also use the group IP address to access the group for management purposes, unless you configured a dedicated management network.
- If you are using IPv6 addresses exclusively in the group, you must use the CLI to change the group IP address. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.
- If you have replication partners configured, you must update those groups with the changes.

Before modifying the group name or group IP address, make sure you understand how these changes will affect your environment:

- You identify replication partners by group name and use the group IP address to perform replication. If you modify the group name or IP address, make sure replication partner administrators make the change to their partner configuration. Replication fails if the partner information is incorrect.
- You use the group IP address as the iSCSI discovery address when connecting initiators to iSCSI targets in the group. If you modify the group IP address, you might need to change your initiator configuration to use the new discovery address.
- You use the group IP address to access the group for management purposes, unless you configured a dedicated management network. If you modify the group IP address, make sure administrators are aware of the change.
- Changing the group IP address disconnects any iSCSI connections to the group and any administrators logged in to the group through the group IP address.
- If you change the IP address of a group for which you are running the GUI locally, or configure a management network for the group, you must uninstall the standalone GUI application and then install it again.
- Applications such as SAN Headquarters use the group IP address to access the group for monitoring purposes. If you modify the group IP address, you must modify the SAN Headquarters configuration to use the new group IP address.
- If you modify the group network configuration (for example, if you modify the group IP address), you might need to make a reciprocal adjustment to the NAS cluster SAN network configuration.

Modify the Group IP Address or Group Name

1. Click **Group → Group Configuration**.
2. Click the **General** tab.
3. Type a new group name.
4. Type a new group IP address.
5. (Optional) Change the location and group description.
6. Click **Save all changes**.

Add a Member to an Existing Group

To add a member to an existing group, you can use either the Remote Setup Wizard for Windows or Linux, or the CLI **setup** command to run the Group Manager setup utility.

The setup utility configures an array as either the first member of a new group or as an additional member of an existing group. During the configuration, the utility prompts you for information such as the name of the member, the name and IP address of the group, and a password for managing group membership. The utility also provides information about what to do after the member has been added. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

When you add a member to a group, keep the following considerations in mind:

- You need the group membership password that is initially established when you create the group. (Note that the group membership password is not the same as the grpadmin password.)
- The new member and the existing group must use compatible firmware levels. Typically, being compatible means having the same major firmware version, but not always. You might need to initially deploy the new member as its own group to have its firmware version be compatible with the existing group.
- The new member must use the same protocol as the group. If the member is joining an IPv6 group, specify the `ipv6` option with the command.
- If you attempt to add a member that does not support IPsec to an existing group in which IPsec is used, **setup** will not be able to add the member to the group.
- If you are adding a member to a group that uses data center bridging (DCB), you must specify the DCB VLAN ID for the new or existing group.

After you add the member, perform the same post-setup tasks as when you initially configured the group. For example:

- Enable the other network ports (see the *Installation and Setup Guide* for your array)

- Enable the management port, if needed (see Configure a Management Network)
- Select the RAID policy (see Set the RAID Policy and Pool for a New Member)

After you complete these tasks, you can configure the new member either through the CLI or through the GUI. To configure the member through the GUI:

1. Click **Group**, expand **Members**, and select the unconfigured member.
2. In the Warning dialog box, click **Yes** or **No**.

   Clicking **Yes** opens the **Configure member** dialog box, in which you select the RAID policy and specify other settings for the member as needed. Click the question mark (**?**) to see more information about these settings.

   Clicking **No** closes the Warning dialog box. When you are ready to configure the RAID policy, you can click the **Configure member** link In the Activities panel to open the dialog box.

NOTE: Before updating a member, verify that it has a default gateway. If it does not, you will receive the following update error: `FTP:Error:Unknown.` If you receive this error, add the default gateway and then you can update the member from the GUI.

Set the RAID Policy and Pool for a New Member

After you add a member to a PS Series group, you must set the RAID policy for the member and choose the storage pool. The storage in the member is available after you set the RAID policy.

Prerequisites

- Review the best practices about setting member RAID policies. For more information about RAID policies on PS Series systems, review the Dell Technical Report entitled *PS Series Storage Arrays: Choosing a Member RAID Policy*, which you can download from the following location:

  http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19861480
- If you used the Remote Setup Wizard to create a group and add the first member to the group, you already set the RAID policy for the member, and the group automatically assigned the member to the default pool. See the *Installation and Setup Guide* for your array.
- Member storage space is available immediately after you have set a RAID policy, although performance is not optimal until the system finishes verifying the member's new RAID configuration.

Procedure

1. On the group to which you want to add the new member, click **Group**.
2. Expand **Members** and then select the member name. The summary window shows whether a member is configured or not.
3. If the member is not configured, a warning is displayed asking if you want to configure RAID now. Click **yes**.
4. In the Activities panel, click **Configure member**.
5. Select the RAID policy in the **Configure member – RAID configuration** dialog box.

   NOTE: You cannot select RAID 50 for PS6610 arrays from the GUI. Instead, you must use the CLI command member select raid-policy. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.
6. (Optional) Select **Wait until the member storage initialization completes**.
7. Click **Next**.
8. Click **Finish** in the Configure Member – Summary dialog box.

NOTE: After initial RAID configuration, it takes a few minutes for Group Manager to display the total usable capacity. Group Manager might show a smaller amount until the process is complete.

Convert a RAID Policy

You configure a RAID policy when you add a member to a group. In most cases, you can convert the RAID policy at a later time.

Prerequisites

- Make sure that you can change the current RAID policy to a different one, and that you understand the conversion options.
- To convert to a no-spare-drives RAID policy, you must use the Group Manager CLI.
- While the RAID policy is changing, the member's RAID status is `expanding`.
- The values in the Member Capacity table are automatically updated, based on the RAID policy that you select.
- To ensure that member space is not available until the RAID verification completes, select **Wait until the member storage initialization completes**, if this option is not selected already.
- Member performance might be temporarily degraded while the RAID policy conversion is taking place.

Procedure

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. Click **Modify RAID configuration** to open the Modify RAID Configuration dialog box.
4. Select the new RAID policy.
5. Click **OK**.

Set RAID Expansion Status

You can pause and resume the start of RAID expansion when new drives are added. This option allows more time when a large number of drives are added.

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. Click **Modify RAID configuration** to open the Modify RAID Configuration dialog box.
4. In the RAID auto expansion panel, select one of the two options:
   - Check **Enable RAID auto expansion** to immediately expand group capacity.
   - Uncheck **Enable RAID auto expansion** to wait until the member storage initialization completes.

   NOTE: You must enable RAID auto expansion after all the drives have been installed.
5. Click **OK**.

Supported RAID Policy Conversions

While a member remains online, you can convert it from one RAID policy to another, but only if the new RAID policy does not reduce the amount of available disk space. For example, converting from RAID 10 to RAID 6 is allowed because it provides more disk space. Converting from RAID 50 to RAID 6 is allowed because it keeps the amount of disk space the same.

Table 24. Supported RAID Policy Conversions shows the supported conversions for each RAID policy.

Table 24. Supported RAID Policy Conversions

Current RAID Policy	Supported Conversion
RAID 10	All
RAID 50	RAID 5, RAID 6
RAID 5	None
RAID 6	None
RAID 6 Accelerated	None

RAID Sets

The tables below show a logical drive layout when an array is initialized for the first time. The actual physical layout of drives can change and evolve due to maintenance and administrative actions. Spare drives can move as they are used to replace failed drives and newly added drives become the spares.

NOTE: It is not possible to determine which physical drives are associated with each RAID set. This information is dynamic and maintained by the Dell EqualLogic firmware.

Table 25. 12–Drive Configuration shows the RAID set relationship for each RAID type in a 12-drive configuration.

Table 25. 12–Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	1	(9+2)	Yes
RAID 10	2	(5+5)	Yes
RAID 50	2	(4+1, 4+1)	For selected configurations
RAID 5	1	(10+1)	Not for business-critical data

Table 26. 14–Drive Configuration shows the RAID set relationship for each RAID type in a 14-drive configuration.

Table 26. 14–Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	1	(11+2)	Yes
RAID 6 Accelerated	1 HDD	(6+2 HDD) (6+2 SSD)	N/A
RAID 10	2	(6+6)	Yes
RAID 50	2	(5+1)(5+1)	For selected configurations
RAID 5	1	(12+1)	Not for business-critical data

Table 27. 16-Drive Configuration shows the RAID set relationship for each RAID type in a 16-drive configuration.

Table 27. 16-Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	1	(13+2)	Yes
RAID 10	2	(7+7)	Yes
RAID 50	2	(6+1, 6+1)	For selected configurations
RAID 5	1	(14+1)	Not for business-critical data

Table 28. 24-Drive Configuration shows the RAID set relationship for each RAID type in a 24-drive configuration.

Table 28. 24-Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	1	(10+2) (9+2)	Yes
RAID 10	2	(6+6) (5+5)	Yes
RAID 50	2	(5+1, 5+1) (4+1, 4+1)	For selected configurations
RAID 5	1	(12+1) (9+1)	Not for business-critical data

Table 29. 42–Drive Configuration shows the RAID set relationship for each RAID type in a 42-drive configuration.

Table 29. 42–Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	2	(12+2) (12+2) (10+2)	Yes
RAID 10	2	(7+7) (7+7) (5+5)	Yes
RAID 50	2	(12+2) (12+2) (10+2)	Not for business-critical data
RAID 5	2	(12+2) (12+2) (10+2)	Not for business-critical data

Table 30. 48-Drive Configuration shows the RAID set relationship for each RAID type in a 48-drive configuration.

Table 30. 48-Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	1	(12+2, 12+2, 12+2) (3+2)	Yes
RAID 10	2	(7+7, 7+7, 7+7) (2+2)	Yes
RAID 50	2	(6+1 ,6+,1, 6+1, 6+1, 6+1, 6+1) (3+1)	For selected configurations
RAID 5	2	(12+1, 12+1, 12+1) (6+1)	Not for business-critical data

Table 31. 84–Drive Configuration shows the RAID set relationship for each RAID type in a 84-drive configuration.

Table 31. 84–Drive Configuration

RAID Policy	Spare Disks	RAID Set Relationship	Best Practices
RAID 6	2	(12+2) (12+2) (12+2) (12+2) (12+2) (10+2)	Yes
RAID 6 Accelerated	1 HHD 1 SSD	(12+2 HDD) (12+2 HDD) (12+2 HDD) (12+2 HDD) (10+2 HDD) (12+2 SSD)	N/A
RAID 10	2	(7+7) (7+7) (7+7) (7+7) (7+7) (6+6)	Yes
RAID 50	2	(12+2) (12+2) (12+2) (12+2) (12+2) (10+2)	Not for business-critical data
RAID 5	2	(12+2) (12+2) (12+2) (12+2) (12+2) (10+2)	Not for business-critical data

Enable and Disable a Volume RAID Preference

A PS Series group uses automatic performance load balancing (enabled by default) to identify the RAID level that provides the best performance for a volume and store volume data on pool members with that RAID level, if such members are available.

Prerequisites

The following prerequisites and considerations apply:

- RAID 5 does not offer optimal data protection for business-critical data due to the higher risks of encountering a second drive failure during a rebuild, and is therefore not recommended. You cannot configure RAID 5 using the Group Manager GUI; however, it can be configured using the CLI.
- Thin clones inherit the RAID preference, if any, of the template volume. You cannot set a separate RAID preference for a thin clone.
- You can override automatic performance load balancing by enabling a RAID level preference on a volume.
- If you enable a RAID preference, the group attempts to store volume data on pool members with that RAID level. The group still uses capacity-based load balancing on the volume.

- To enable a volume RAID preference, make sure that at least one member in the volume's pool has the preferred RAID level. If no pool member has the preferred RAID level, the group ignores the RAID preference until a member exists with the preferred RAID level.
- If you disable a RAID preference on a volume, the group resumes automatic performance load balancing.

Procedure

To enable or disable a volume RAID preference:

1. Click **Volumes**.
2. Expand **Volumes** and select the volume name.
3. Click **Modify settings** and then click the **Advanced** tab to open the Modify volume settings – Advanced dialog box.
4. Select either:
   - Preferred RAID level under **Volume RAID preference**.
   - **Automatic** to disable a RAID preference
5. Click **OK**.

About Overriding Automatic Load Balancing

A PS Series group uses automatic performance load balancing (enabled by default) to identify the RAID level that provides the best performance for a volume and store volume data on pool members with that RAID level, if such members are available.

You can override automatic performance load balancing (or ignore any RAID preference for the volume) by binding a volume to a specific pool member.

If you bind a volume to a pool member, the group stores the volume data on the member, instead of distributing data across multiple pool members.

The following restrictions apply:

- You can bind a volume only to a member that is in the same pool as the volume.
- If you bind a volume to a member and then delete that member from the pool or group, the group cancels the bind operation.
- Thin clones inherit the member binding setting, if any, of the template volume. You cannot have a separate member binding setting for a thin clone.
- You cannot enable synchronous replication on a volume that is bound to a group member.
- You cannot use the Group Manager GUI to bind a volume to a member. Instead, you must use the following Group Manager CLI command line:

  ```
  volume select volume_name bind member_name
  ```

  To unbind a volume from a member, use the following CLI command line:

  ```
  volume select volume_name unbind
  ```

See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about using CLI commands.

Shut Down a Group

To perform maintenance, you might need to shut down a PS Series group. While the group is shut down, volumes are not available.

1. Stop all I/O applications connected to the group.
2. Disconnect iSCSI initiators from the group.
3. Shut down each group member. Do not turn off power to the PS Series array until you cleanly shut down the member.
   To shut down a member:
   a. Click **Group**.

   b. Expand **Members** and select the member that you want to shut down.
   c. Click the **Maintenance** tab.
   d. In the Power panel, click the **Shut down** button.
   e. In the **Member shutdown** dialog box, enter the grpadmin password and click **OK**. (Only group administrators can shut down members.)
   f. The system displays a warning message that the member is shutting down. Click **OK** to acknowledge the warning.

4. To turn off array power, turn off all power switches on the array after the shutdown completes.

To restart the group, power on all group members.

Create an Empty Storage Pool

1. Click **Group**.
2. Select **Storage Pools**.
3. Click **Create storage pool** to open the Create Storage Pool dialog box.
4. Specify a pool name and, optionally, a description.
   - Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - The pool description can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. Click **Next** → **Finish**.

   NOTE: Alternatively, you can create a storage pool from an existing member. For instructions, see Create a Storage Pool from an Existing Member.

Create a Storage Pool from an Existing Member

NOTE: This procedure should be treated as an alternative means for creating a storage pool. See Create an Empty Storage Pool for more information.

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. In the Activities panel, click **Move member**.
4. In the dialog box, click **Create new pool** to open the Create Storage Pool dialog box.
5. Type a name and description for the new pool.
   - Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - The pool description can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
6. Click **Finish**.
7. Click **OK** in the Move Member dialog box.
8. Confirm that you want to create the pool.

The member status shows as `moving` until the move operation completes.

When you move a member from one pool to another, the volumes using space on that member stay in the original pool. Therefore, the other members of the original pool must contain enough space to store any volume data that was on the member being moved. Moving a member to a different pool can take a long time, depending on the amount of data that the group must move.

If necessary, you can cancel an in-progress member pool move operation. The member immediately returns to the original pool.

To cancel an in-progress member pool move operation:

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. Click **Cancel member move**.

Change a Storage Pool Name or Description

You can change the name or description of any storage pool, including the default pool.

1. Click **Group**.
2. Expand **Storage Pools** and then select the pool name.
3. Click **Modify pool settings**.
4. Modify the pool name or description.
   - Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - The pool description can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. Click **OK**.

Merge Storage Pools

You can merge any storage pool except the default pool into another pool, called the destination pool. Merging pools moves the pool members and volume data into the destination pool. The group then deletes the empty pool.

To merge storage pools:

1. Click **Group**.
2. Expand **Storage pools** and then select the pool to be merged with another pool.
3. Click **Merge storage pool** to open the Merge Storage Pools dialog box.
4. Select the destination pool and click **OK**.

Delete a Storage Pool

When you delete a populated storage pool, the group immediately moves its members and volumes to the default pool.

NOTE: You cannot delete the default storage pool, but you can rename it.

1. Click **Group**.
2. Expand **Storage Pools** and then select the pool name.
3. Click **Delete storage pool**.
4. Confirm the pool deletion.

About Groupwide Volume Defaults

The groupwide volume default settings control certain attributes of all volumes that you create on the group, including the snapshot space, snapshot behavior, thin-provisioning space, sector size, and iSCSI alias naming. You can override many of these settings on a per-volume basis when you create a volume, but unless otherwise specified, these settings are the defaults for all volumes on the group.

You can change the default values to meet the needs of your environment.

Modify Groupwide Volume Settings

When you create or enable thin provisioning on a volume, the group applies defaults unless you explicitly override them for a volume. These defaults control snapshot space, snapshot behavior, thin-provisioning space, sector size, and iSCSI alias naming.

You can modify the groupwide default values to meet the needs of your configuration.

NOTE: Changes to the groupwide default values apply only to new volumes.

To modify the settings:

1. Click **Group** → **Group Configuration**.
2. Click **Defaults** to open the Settings window.
3. Change the default settings as needed.
4. Save any changes that you made.

About Space Borrowing

Space borrowing includes snapshot space borrowing and replication borrowing.

(For more information about these borrowing methods, see About Snapshot Space Borrowing and About Replication Borrowing.)

PS Series firmware v8.0 and later provides the following borrowing capabilities:

- Extends the existing snapshot and local replica borrowing capability to enable borrowing from unused delegated space
- Provides borrowing to remote replicas so that they can borrow space beyond their replica reserve (similar to snapshot borrowing)

Space is potentially available for borrowing from the following areas:

- Unused snapshot and replication reserves
- Unused delegated space
- Free space

You should be aware of the following restrictions regarding space borrowing:

- Borrowed space is intended to help during peaks of activity when more space is needed temporarily. It does not take the place of carefully provisioning reserves.
- Remote replicas can borrow beyond their associated remote replica reserve, but the total amount of reserve space must still fit within the delegated space.
- When necessary, the system will automatically delete an object that is borrowing space. Any objects in borrowed space are considered to be “at risk.”

Benefits of Space Borrowing

Space borrowing provides the following benefits:

- Eases the effects of setting reserves either too high or too low, and helps alleviate the pressure of predicting reserve usage accurately, especially if this usage changes frequently.
- Allows snapshots and replicas to borrow enough space to hold the copies that would otherwise be deleted. As long as the pool has enough unused space to borrow, the copy can borrow space to satisfy the max-keep policy.

The borrowing capabilities added for v8.0 and later can help in the following situations:

- If a volume’s snapshot reserve or replica reserve is set too high or too low:
  - Too high – The system considers the unused snapshot and replica reserve as “borrowable” space, which is then available to other volumes whose reserves were set too low. This flexibility accommodates volumes whose IO patterns caused the volume reserves to be exceeded and better utilizes the available space.
  - Too low – The system allows snapshots and replicas to borrow space, which maintains the max-keep policy. Older snapshots and replicas will not need to be deleted if the system can borrow space for them.
- If replica reserve is configured to be too small or too large:
  - Too small – The system allows the replicas to borrow from outside their replica reserve, which maintains the max-keep policy.
  - Too large – The system enables other snapshots and replica sets in the pool to borrow the unused space.

Displaying Space-Borrowing Information

The Group Manager GUI can help you monitor borrowed space on your system. To see borrowing information, either:

- Click **Group** → **Borrowed Space**.

> **NOTE: This display applies only to the space in the pool that is available for borrowing, not to *all* of the space in the pool. Volume reserve and storage containers do not appear in the tables because that space cannot be borrowed.**

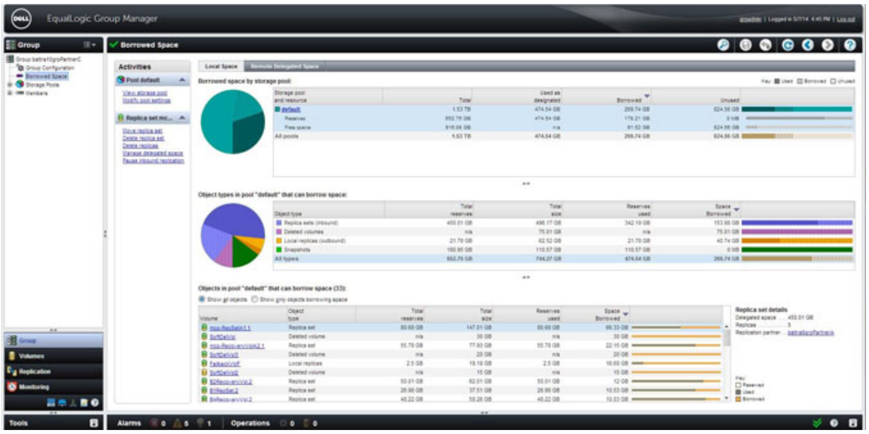

Borrowed space by storage pool:

Storage pool and resource	Total	Used as designated	Borrowed	Unused
default	1.82 TB	474.54 GB	269.74 GB	824.56 GB
Reserves	652.75 GB	474.54 GB	178.21 GB	0 MB
Free space	916.08 GB	n/a	91.52 GB	824.56 GB
All pools	1.82 TB	474.54 GB	269.74 GB	824.56 GB

Object types in pool "default" that can borrow space:

Object type	Total reserves	Total size	Reserves used	Space Borrowed
Replica sets (inbound)	450.01 GB	496.17 GB	342.19 GB	153.98 GB
Deleted volumes	n/a	75.01 GB	n/a	75.01 GB
Local replicas (outbound)	21.78 GB	62.52 GB	21.78 GB	40.74 GB
Snapshots	180.95 GB	110.57 GB	110.57 GB	0 MB
All types	652.75 GB	744.27 GB	474.54 GB	269.74 GB

Objects in pool "default" that can borrow space (33):

● Show all objects ○ Show only objects borrowing space

Volume	Object type	Total reserves	Total size	Reserves used	Space Borrowed
mcp-RecSetA1.1	Replica set	80.68 GB	147.01 GB	80.68 GB	66.33 GB
SoftDelVol	Deleted volume	n/a	30 GB	n/a	30 GB
mcp-RecoveryVolA2.1	Replica set	55.78 GB	77.93 GB	55.78 GB	22.15 GB
SoftDelVol3	Deleted volume	n/a	20 GB	n/a	20 GB
FailbackVolF	Local replicas	2.5 GB	19.19 GB	2.5 GB	16.68 GB
SoftDelVol2	Deleted volume	n/a	15 GB	n/a	15 GB
B2RecoveryVol2	Replica set	50.01 GB	62.01 GB	50.01 GB	12 GB
B1RecSet2	Replica set	26.98 GB	37.51 GB	26.98 GB	10.53 GB
B4RecoveryVol2	Replica set	40.22 GB	58.26 GB	40.22 GB	18.03 GB

Replica set details
Delegated space ..... 450.01 GB
Replicas 5
Replication partner ... battra6grpPartnerA

Key: □ Reserved ■ Used ■ Borrowed

- Click the **Total borrowed space** link from either the Group Disk Space panel or the Pool Disk Space panel.

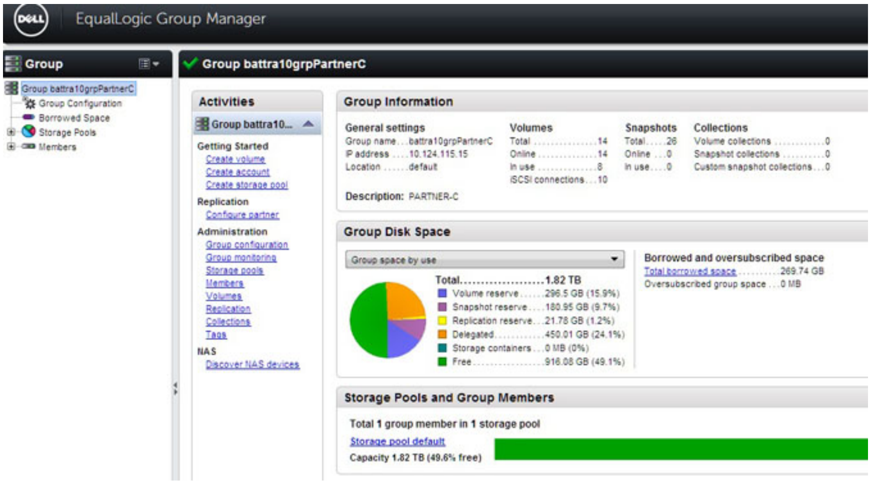

Group battra10grpPartnerC

Group Information

General settings
Group name ... battra10grpPartnerC
IP address 10.124.115.15
Location ...... default

Volumes
Total 14
Online 14
In use 8
iSCSI connections ... 10

Snapshots
Total 26
Online ... 0
In use 0

Collections
Volume collections 0
Snapshot collections 0
Custom snapshot collections ... 0

Description: PARTNER-C

Group Disk Space

Group space by use

Total..................... 1.82 TB
Volume reserve ...... 296.5 GB (15.9%)
Snapshot reserve .... 180.95 GB (9.7%)
Replication reserve ... 21.78 GB (1.2%)
Delegated .... 450.01 GB (24.1%)
Storage containers ... 0 MB (0%)
Free .... 916.08 GB (49.1%)

Borrowed and oversubscribed space
Total borrowed space .... 269.74 GB
Oversubscribed group space ... 0 MB

Storage Pools and Group Members

Total 1 group member in 1 storage pool
Storage pool default
Capacity 1.82 TB (49.6% free)

You can also see these statistics through several CLI commands. Refer to the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

Using Space-Borrowing Information

Using information displayed in the GUI, you can see how space borrowing is affecting your system and answer the following questions:

1. Have I set my reserves too high or too low?
2. Am I benefiting from the borrowing feature?
3. Are any objects at risk of being deleted because they are borrowing space?
4. How much space can I use without causing objects within borrowed space to be deleted?

The information on the Borrowed Space screen can help answer these questions. See the following table.

Question	Answer
Reserves too high or too low?	To determine if you have set the reserves too high or too low, you can see whether or not any object is borrowing and also the extent of the borrowing. For each volume or replica set, you can see the amount of space reserved and the amount of space in use. The amount of in-use space includes the amount of space that the snapshot or replica is borrowing. To determine whether the reserves are set appropriately, you should continually monitor the volume's snapshot or replica borrowing.
Any benefits from borrowing?	For objects that are borrowing space, you can monitor the number of snapshots or replicas to determine whether the correct number of them are being retained. Various GUI displays and CLI commands enable you to see this information. The GUI displays indicate on a per-pool basis whether borrowing is occurring and the extent of the borrowing. Additionally, you can see if a particular object is borrowing and whether replica sets on a partner are borrowing.
Any objects at risk of being deleted?	You can determine if any objects are at risk of being deleted by monitoring whether an object is borrowing space. If no borrowing is occurring, the object is not at risk.
Any objects at risk of being deleted as the result of using too much space?	To determine whether an object is at risk of being deleted, you can compare the amount of a resource that is being borrowed (such as snapshot reserve, replication reserve, or free pool space) to the amount that is not being used.

About Compression of Snapshots and Replicas

The compression feature maximizes available space within the storage pool by automatically targeting inactive snapshots and replicas and compressing them to reduce their storage footprint. When compression is enabled, the array periodically monitors activity to determine which snapshots and replicas can be safely compressed. Compressed data is automatically "rehydrated" (that is, decompressed) when that information needs to be accessed by the array.

Compression is disabled by default on Dell EqualLogic arrays. If compression is enabled, it must remain enabled to provide access to already compressed data. Compression can be manually suspended and resumed, as needed, to prevent the array from compressing new snapshots and replicas, but any previously compressed data remains in its compressed state.

When a compression-enabled member is part of a group containing other (non-compression) members, compressed data might be off-loaded to those other members, but this occurs only if the compression-enabled member becomes full of compressed data.

When data is moved from one array to another, such as during a "vacate" action, the compressed data must be rehydrated first. After the data is moved, it is again eligible for compression if the new destination supports it. Compression can be enabled on any PS6210 or PS6610 array that is running compatible firmware.

Compression Prerequisites

The compression feature is available only on PS6210 and PS6610 arrays running firmware version 8.0 or later.

Compression is disabled by default on all arrays until it is manually activated either through the Group Manager interface or the command-line interface (CLI).

NOTE: The array monitors the storage environment for eligible snapshots. After starting compression, it might be several hours before any compression activity is displayed.

About Rehydration

Rehydration is the process whereby data that has been compressed becomes decompressed and readily accessible. Rehydration happens automatically and occurs as necessary to provide access to compressed data. Whenever data is vacated from one group member to another, any compressed data is automatically rehydrated during that process. This rehydration occurs regardless of whether the receiving member is compression-capable or not.

About Compression Statistics

Compression statistics assess the overall space-savings achieved within a storage pool, a member, or a volume, including replicated volumes. These statistics are accessible through the GUI or the CLI. After starting compression initially or resuming it from a suspended state, it might be several hours before any new compression activity occurs.

NOTE: When a compressed snapshot is cloned or rolled back, the compression statistics will indicate that the space has been rehydrated. Base volumes with compressed pages are not supported.

Compression Statistics by Pool

Pool-level compression statistics allow you to assess the overall space-savings achieved by all compression-capable members in a storage pool. The following status information is provided:

- **Expanded size –** The space required if all the compressed snapshots stored on all members of the storage pool were rehydrated.
- **Compressed size –** The space currently being used by all compressed snapshots on all members of the pool.
- **Space savings –** The amount of space being saved in the storage pool; the difference between the expanded size and the compressed size.

Compression Statistics by Member

Member-level compression statistics allow you to assess the overall space-savings achieved within a member as a result of using the compression feature. The following status information is provided:

- **Expanded size –** The space required if all compressed snapshots stored on that member were rehydrated.
- **Compressed size –** The space currently being used by all compressed snapshots on that member.
- **Space savings –** The amount of space being saved on the member; the difference between the expanded size and the compressed size.

Compression Statistics by Volume

Volume-level compression statistics allow you to assess the overall space-savings achieved within a volume as a result of using the compression feature. The following status information is provided:

- **Expanded size –** The space required if all compressed snapshots in that volume were rehydrated.
- **Compressed size –** The space currently being used by all compressed snapshots in that volume.
- **Space savings –** The amount of space being saved on the volume; the difference between the expanded size and the compressed size.

Member Compression States

Table 32. Member Compression States describes the possible member compression states.

Table 32. Member Compression States

Member State	Description
No-Capable-Hardware	Compression cannot be enabled on this member due to its hardware capabilities.
Not Started	Compression has never been successfully enabled on this member.
Running	Compression is currently running on this member.
Suspended	Compression has previously been enabled and is now paused.
Unknown	The compression state of the member cannot be determined due to the member being offline or an internal error condition. **NOTE: This status is visible only within the pool area.**

Enable Compression

Snapshot and replica compression can be enabled on any PS6210 or PS6610 array running the supported firmware. Compression is enabled at the pool level. If the group contains more than one compression-enabled member, compression must be enabled on each member individually.

1. Click **Group**.
2. Expand **Members** and select the group member on which compression is to be enabled.
3. In the Activities panel, click **Start compression**.

   The Compression Settings dialog box opens and reminds that compression cannot be undone but only suspended after it is enabled.
4. Click the **Continue** button.

The General Member Information panel shows `Compression ... running` under the General settings section, and the Activities panel contains a **Suspend compression** link, which is used to temporarily pause compression.

NOTE: When compression is initially enabled or resumed after being suspended, the display for compression statistics might not change for several hours as the array monitors activity.

Suspend Compression

Snapshot and replica compression can be temporarily suspended on a member to immediately stop the compression of new snapshots. This process does not rehydrate already compressed data; all compressed data remains in its compressed state when compression is suspended.

1. Click **Group**.
2. Expand **Members** and select the group member on which compression is to be suspended.
3. In the Activities panel, click **Suspend compression**.

The General Member Information panel shows `Compression ... suspended` under the General settings section, and the Activities panel contains a **Resume compression** link, which is used to reactivate compression.

NOTE: When compression is initially enabled or resumed after being suspended, the compression statistics might not change for several hours as the array monitors activity.

Resume Compression

To resume compression on a member after it has been suspended:

1. Click **Group**.
2. Expand **Members** and select the group member on which compression is to be resumed.
3. In the Activities panel, click **Resume compression**.

The General Member Information panel shows `Compression ... running` under the General settings section, and the Activities panel again contains the **Suspend compression** link.

NOTE: When compression is initially enabled or resumed after being suspended, the display for compression statistics might not change for several hours as the array monitors activity.

View Compression Statistics by Pool

You can view the compression statistics of an entire storage pool from any compression-enabled member.

1. Click **Group**.
2. Expand **Storage Pools** and select the pool for which you want to view compression statistics.
3. Click the **Status** tab and locate the Pool Disk Space panel.

   The statistics are listed under **Snapshot and replica set compression**.

See Compression Statistics by Pool.

View Compression Statistics by Member

You can view compression statistics for any compression-enabled member.

1. Click **Group**.
2. Expand **Members** and select the member for which you want to view compression statistics.
3. Click the **Status** tab and locate the Member Space panel.

   The statistics are listed under **Snapshot and replica set compression**.

See Compression Statistics by Member.

View Compression Statistics by Volume

You can view the compression statistics of specific volumes from any compression-enabled member.

1. Click **Volumes**.
2. Expand **Volumes** and select the volume for which you want to view compression statistics.
3. Click the **Status** tab and locate the Volume and Snapshot Space panel.

   The statistics are listed under **Snapshot compression**.

You can view the compression statistics of replicated volumes as well from any compression-enabled member that has been configured as a replication partner.

1. Click **Replication**.
2. Expand the replication partner containing the volume that you want to view.
3. Click **Volume Replication**.
   The compression statistics can be found in the Volume Replication panels.

See Compression Statistics by Volume.

Compression Commands in the CLI

Snapshot and replica compression can be enabled using the command-line interface (CLI) on any PS6210 or PS6610 array running firmware version 8.0 or later.

Table 33. CLI Commands for Compression

Action	Command
Start compression	member select ***member_name*** data-reduction compression start
Pause compression	member select ***member_name*** data-reduction compression suspend
Resume compression	member select ***member_name*** data-reduction compression resume
Display member compression statistics	member select ***member_name*** show
Display volume compression statistics	volume select ***volume_name*** show
Display pool compression statistics	pool select ***pool_name*** show

The General settings section of the General Member Information panel shows `Compression ... running`, and the Activities panel contains the **Suspend compression** link, which is used to temporarily pause compression.

NOTE: When compression is initially enabled or resumed after being suspended, the display for compression statistics might not change for several hours as the array monitors activity.

9

About Volumes

Volumes provide the storage allocation structure within the PS Series group.

To access storage in a PS Series group, you allocate portions of a storage pool to volumes. You can create a volume on a single group member or one that spans multiple group members. You assign each volume a name, size, and a storage pool. The group automatically load balances volume data across pool members.

Create a Volume

To create a new volume:

1. Click **Volumes → Volumes**.
2. In the Activities panel, click **Create volume**.
3. Follow the steps in the **Create volume** wizard to specify volume settings, space settings, iSCSI access points, access types, and sector size.
4. Click **Finish**.

Modify a Volume Name or Description

You can modify a volume name or description any time after you create a volume, subject to the following considerations and constraints:

- If you modify a volume name, the iSCSI target name (and any snapshot or replica set names) does not change. However, if you modify a volume name, and the volume alias is set to be the same as the volume name, the alias also changes.
- If you modify the name of a replicated volume, you continue to identify the replica set on the secondary group by the original volume name.
- Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

To modify volume information:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **General** tab to open the Modify Volume Settings – General dialog box.
5. Modify the name and, optionally, the description.
6. Click **OK**.

Modify a Volume Permission

A volume can have read-write or read-only permission, unless it is a template volume.

The following considerations apply:

- To change a volume permission to read-only, you must first set the volume offline.

- You cannot set a template volume to read-write permission.

To modify a volume permission:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Set access type**.
4. Change the permission in the **Set access type** dialog box.
5. (Optional) Select the **Allow simultaneous connections from initiators with different IQNs** checkbox.
6. Click **OK**.

Modify a Volume Alias

An alias can help administrators identify a volume. For example, some iSCSI initiators display the volume alias in addition to the iSCSI target name.

The following constraints apply:

- When you create a volume, it has an alias only if the groupwide default is to use the volume name as the alias. Otherwise, the volume does not have an alias.
- A volume alias is a name. Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

To modify a volume alias:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **Advanced** tab to open the Modify Volume Settings – Advanced dialog box.
5. Specify a volume alias in the **Public alias** field.
6. Click **OK**.

Modify the Administrator for a Volume

You can modify the administrator for a volume only if the group has more than one group administrator (grpadmin).

To create another administrator for a group, see Create a Local Administration Account.

To control which administrators have access to a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **General** tab to open the Modify Volume Settings – General dialog box.
5. In the **Volume administrator** field, change the name of the administrator or select **none**.
6. Click **OK**.

About Smart Tags

Smart tags provide a mechanism for PS group administrators to organize, search, and filter the volumes in their groups. Volumes can be tagged with new or predefined tags, and a volume can have more than one tag associated to it. This feature makes it easier for the group administrator to manage a group with a large number of volumes.

Smart tags can be either simple tags or extended tags:

- A simple tag has no values. For example, the tag “Backup” indicates that the volume is used as a backup volume without going into further detail.
- An extended tag uses a main tag and a value. For example, the tag “Applications” with the value “Sharepoint” indicates that the volume information is used within that application.

Within Group Manager, only group administrators are allowed to create, rename, and delete tags. However, administrators with write privileges for the volume can assign tags to the volume. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

Predefined Smart Tags

Smart tags include a set of predefined tags, described in Table 34. Predefined Smart Tags. These tags can be edited and deleted, but they cannot be restored to their original settings.

Table 34. Predefined Smart Tags

Predefined Tag	Description
Applications	Tag the volume based off the application that you define the volume with. The application tag name has predefined values to select from: • Database-Server • File-Server • FTP • Mail-Server • Sharepoint
Virtualization	Associates the volume with a virtualization tag name, either Yes or No
Backup	Associates the volume with a backup tag name, either Yes or No
Desktops	Associates the volume with a desktops tag name, either Yes or No
Location	Associates the volume with a specific location. The location tag name has predefined values to select from: • Boston • New-York • London • Shanghai

Using Smart Tags

With smart tags, you can:

- Associate tags and tag values for all the volumes in a group
- Associate tags and tag values for a specific volume in a group
- Filter the information displayed for volumes

To associate tags for all volumes in a group, you can use either the Group pane or the Volumes pane.

From the Group pane:

1. Click **Group** → **Group *group_name***.
2. In the Activities panel, click **Tags** to open the Manage Tags dialog box.

From the Volumes pane:

1. Click **Volumes** → **Volumes**.
2. In the Activities panel, click **Manage tags** to open the Manage Tags dialog box.

To associate tags for a specific volume, you can use either the Volumes panel, the Activities panel, or the General Volume Information panel.

From the Volumes panel:

1. Click **Volumes** → **Volumes**.
2. Select a volume from the list displayed in the Volumes panel.
3. In the Activities panel, click **Modify tags** to open the Pick Tags for Volume dialog box.

From the Activities panel:

1. Click **Volumes**.
2. Expand **Volumes** and select a volume from the tree view.
3. In the Activities panel, click **Modify tags** to open the Pick Tags for Volume dialog box.

From the General Volume Information panel:

1. Click **Volumes**.
2. Expand **Volumes** and select a volume from the tree view.
3. In the General Volume Information panel, click **Tags** to open the Pick Tags for Volume dialog box.

To filter the information displayed in the Volumes panel, you can use either tags or column headings.

To filter by tags:

1. Click **Volumes** → **Volumes**.
2. In the Volumes panel, click **Settings** to open the Filter Volumes dialog box.

To filter by column headings:

1. Click **Volumes** → **Volumes**.
2. In the Volumes panel, click **Pick tag columns** to open the Pick Tags for Columns dialog box.

Set a Volume Offline or Online

By default, when you create a volume, the group sets the volume online. An iSCSI initiator can discover or connect to an online volume only. To make a volume inaccessible to iSCSI initiators, set the volume offline; the group closes all current iSCSI connections to the volume.

NOTE: To set a volume online, each member that contains volume data must be online.

To set a volume online:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Set online**.

To set a volume offline:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Set offline**.
4. Confirm that you want to set the volume offline.

Delete a Volume

When you delete a volume, space that the group allocated to the volume becomes part of free pool space.

The following requirements and considerations apply:

- If you delete a volume, the group also deletes its snapshots. However, the group does not delete any volume replicas on the secondary group.
- The volume must be set offline to perform the delete operation. The group closes any active iSCSI connections to the volume.
- You cannot delete a template volume if:
  - The volume is attached to thin clones.
  - Recovery thin clones exist.

To delete a volume:

1. Click **Volumes**.
2. Expand and then select the volume name.
3. In the Activities panel, click **Set volume offline**.
4. Select the volume name (if it was deselected).
5. Click one of the following volume types:
   - **Delete volume**
   - **Delete template**
   - **Delete thin clone**
6. Confirm that you want to delete the volume and its data.

About Volume Collections

You can group multiple volumes for the purpose of performing an operation simultaneously on the volumes. Volume collections are useful when you have multiple, related volumes on which you would like to simultaneously perform maintenance tasks.

If you have a set of volumes that are organizationally related to one another, but do not require any operational grouping, use volume folders to organize them.

A volume collection includes one or more volumes from any pool. In a single operation, you can create snapshots of the volumes (a snapshot collection) or replicas of the volumes (a replica collection). You can also enable synchronous replication (SyncRep) for volume collections.

NOTE: You cannot use a template volume in a volume collection.

Create a Volume Collection

To create a volume collection:

1. Click **Volumes** → **Volume Collections**.
2. Click **Create volume collection** to open the Create Volume Collection dialog box.
3. In the **General Settings** section of the dialog box, specify a name and an optional description for the collection.
4. Click **Next** to go to the Components section of the dialog box.
5. Select the volumes to include in the collection.
6. Click **Next** to review the Summary section.
7. Click **Finish**, or click **Back** to make changes.

The volume collection is displayed under **Volume Collections**.

Modify a Volume Collection

To modify a volume collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection.
3. Click **Modify volume collection** to open the Modify Volume Collection dialog box.
4. Click the **General** tab to change the collection name or description.
5. Modify the name (up to 63 characters) or description (up to 127 characters).

   Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
6. Click the **Components** tab to add volumes to, or remove volumes from, the collection.
7. Select and deselect volumes as needed.
8. Click **OK**.

Delete a Volume Collection

You can delete volume collections that you created.

Deleting a volume collection does not delete the volumes in the collection or any snapshots or replicas. However, the group deletes any schedules for the volume collections.

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection.
3. Click **Delete volume collection**.
4. Confirm that you want to delete the collection.

About Volume Folders

With volume folders, you can organize volumes into folders in the Group Manager GUI. These folders provide a quick visual reference.

Folders are an organizational tool only; they do not affect the volumes they contain. If you have a set of volumes on which you would like to simultaneously perform operations, use volume collections.

Unlike volume collections, volume folders display the amount of space that each volume has borrowed for snapshots, as well as the number of iSCSI connections to each volume.

Volume Folder Configuration Considerations

The following considerations apply to configuring volume folders:

- You can move volumes into, out of, or between folders only if your account has group administrator privileges.
- Each group can have a maximum of 1024 volume folders.
- The folders cannot be nested; only one level of folders is permitted.
- Each folder name must be unique. Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- If you have administrator privileges, you can enter a 127-character description of each folder. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

- Deleting a volume folder does not delete the volumes it contains.
- A volume folder cannot contain a volume collection.
- A volume folder can contain standard, thin-provisioned, template, thin clone, or synchronous replication (SyncRep) volumes.
- A volume folder cannot contain failback replica sets, or promoted or cloned inbound replica sets.
- When the last volume is removed from a folder, the folder is not deleted.
- Volume administrators can view all of the volumes in a folder, provided that the volumes were not created by other volume administrators.

Create a Volume Folder

1. Click the **Volumes** tab.
2. In the Activities panel, click **Create volume folder** to open the dialog box.
3. In the dialog box, type a name for the volume folder in the **Name** field.

   Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
4. (Optional) Add a description of the folder in the **Description** field.

   The description can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. Click **OK**.

The new volume folder is displayed in the tree view below the volumes in the group. (If folders are not being displayed, you can enable the option from the tree-view menu (▤ ▾). Select **Group volumes** → **Show volume folders**.

Display Volume Folders

1. Click the **Volumes** tab.
2. Click the tree-view options menu next to the Volumes heading.
3. Select **Group volumes** → **Show volume folders**.

Volume folders are displayed in the tree view below the volumes in the group.

Hide Volume Folders

1. Click the **Volumes** tab.
2. Click the tree-view options menu next to the Volumes heading.
3. Select **Group volumes** and then clear **Show volume folders**.

Volume folders are not displayed in the tree view below the volumes in the group.

Rename Volume Folders

1. Click the **Volumes** tab.
2. Expand **Volumes** and then select the folder that you want to rename.

   If volume folders are not being displayed, you can enable the option from the tree-view menu (▤ ▾). Select **Group volumes** → **Show volume folders**.
3. In the Activities panel, click **Rename folder** to open the dialog box.
4. Change the name of the folder in the **Name** field.

   Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

5. (Optional) You can also change the description of the folder in the **Description** field.

   The description can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
6. Click **OK**. The folder displays its new name.

Add and Remove Volumes from Folders

You can move volumes into, out of, or between folders. Your account must have group administrator privileges.

If volume folders are not being displayed, you can enable the option from the tree view menu (▤ ▾). Select **Group volumes → Show volume folders**.

Add Volumes to a Volume Folder

1. Click the **Volumes** tab.
2. Expand **Volumes** to display a list of available volumes.
3. Select the name of the volume that you want to add to the folder.
4. In the Activities panel, click **Move to folder**. A dialog box opens with a list of the available volume folders.
5. Select the volume folder where you want to add the volume and then click **OK**.

The volume is displayed within the selected volume folder in the tree.

To show the volumes contained within a volume folder, click the name of the folder.

Remove Volumes from a Volume Folder

1. Click the **Volumes** tab.
2. Expand **Volumes** to display a list of available volumes.
3. Select the name of the volume that you want to remove from the folder.
4. In the Activities panel, click **Remove from folder**.
5. Click **Yes**.

The volume is removed from the volume folder and is displayed in the tree view below Volumes.

Move Volumes Between Volume Folders

1. Click the **Volumes** tab.
2. Expand **Volumes**.
3. Select and then expand the name of the volume folder to display the volumes grouped within that folder.

   If volume folders are not being displayed, you can enable the option from the tree-view menu (▤ ▾). Select **Group volumes → Show volume folders**.
4. Select the name of the volume that you want to remove from the folder.
5. In the Activities panel, click **Move to another folder**. A dialog box opens that lists all of the available folders.
6. Select the name of the folder to which you want to move the volume.
7. Click **OK**.

NOTE: You can move volumes into, out of, or between folders. Your account must have group administrator privileges.

Delete Volume Folders

When you delete a volume folder, the volumes contained with the folder are not deleted, but are moved to the Volumes list.

If volume folders are not being displayed, enable the option from the tree view menu (▤ ▾): **Group volumes → Show volume folders**

1. Click the **Volumes** tab.

2. Expand **Volumes** and then select the name of the volume folder that you want to delete.
3. In the Activities panel, click **Delete folder** to open the dialog box.
4. Click **Yes** to delete the folder.

   If the deleted folder contained volumes, the volumes are displayed in the Volumes list in the tree view.

About Restoring Deleted Volumes

Volume undelete provides an administrator with the ability to restore volumes that might have been deleted by mistake. This feature is enabled by default. To turn it off or to turn it back on after it has been disabled, use the CLI recovery-bin volume command.

Volumes are automatically moved to the recovery bin by the firmware whenever a volume is deleted by an administrator. They remain in the recovery bin for 1 week before they are automatically deleted by the firmware. You can also manually purge them from the recovery bin.

When you delete a volume, the firmware automatically converts its volume type to a thin-provisioned volume. Restoring the volume converts it back to its original type, and restores all the user data residing on that volume. If not enough space is available to restore a volume to its original type, the volume is restored as a thin-provisioned volume.

The following types of volumes can be restored after deletion:

- Standard volumes
- Recovery volumes
- Synchronous replication (SyncRep) volumes residing in the SyncActive pool
- Thin-provisioned volumes

NOTE:

- **If you delete a volume for which synchronous replication is enabled, the system will place the SyncActive volume into the recovery bin. However, the SyncAlternate volume will be deleted outright and cannot be recovered.**
- **If the normal use of pool space results in a low amount of available free space, the firmware will permanently delete volumes in the recovery bin to reclaim free space. This deletion can occur before the 1-week time limit has been reached.**
- **Both snapshot space borrowing and the recovery bin make temporary use of free space in a pool. When the amount of free space in a pool becomes low, recoverable volumes will be purged from the recovery bin before snapshots that use borrowed space are deleted.**
- **Manual purge operations are displayed in the audit log. Automatic purge operations performed by the firmware are logged in the event log.**

Enable or Disable Volume Undelete

Volume undelete is enabled by default. Disabling volume undelete immediately purges all of the volumes that are preserved in the recovery bin.

To turn off volume delete or to turn it back on after it has been disabled, use the CLI command recovery-bin. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

Display Deleted Volumes

1. Click **Volumes**.
2. Click **Volumes** in the Volumes panel (not an individual volume name).
3. In the Activities panel, click **Manage recovery bin**. The Volume Recovery Bin dialog box opens.

   - The recovery bin lists the number of deleted volumes that can be recovered, as well as each volume name, its original storage pool, and size. The recovery bin also shows the date and time when the volume was deleted and the time when the volume will automatically be purged if it has not been restored before then.
   - Volumes that contain no data are not preserved in the recovery bin. When you delete a volume that has no data, it is deleted immediately.

- Deleted volumes remain in the recovery bin for up to 1 week after deletion. If a deleted volume has not been restored after 1 week, it will be purged after the date and time shown in the Volume Recovery Bin dialog box.
- When a volume has been deleted, its information appears slightly differently in the CLI than in the GUI. Whereas the GUI shows the original name of the volume even after it has been deleted, the CLI shows a modified name for the volume when you list the contents of the recovery bin.

NOTE:

- **If the amount of available free space within the group decreases to less than the specified system level, the firmware will automatically begin to purge deleted volumes, even if a week has not elapsed.**
- **If you delete a volume for which synchronous replication is enabled, the system will place the SyncActive volume into the recovery bin. However, the SyncAlternate volume will be deleted outright and cannot be recovered.**

Restore Deleted Volumes

When you delete a volume, the firmware automatically converts its volume type to a thin-provisioned volume. Restoring (undeleting) the volume converts it back to its original type, and restores all the user data residing on that volume. If not enough space is available to restore a volume to its original type, it is restored as a thin-provisioned volume.

To restore a deleted volume:

1. Click **Volumes**.
2. Click **Volumes** in the Volumes panel (not an individual volume name).
3. In the Activities panel, click **Manage recovery bin**. The Volume Recovery Bin dialog box opens.
4. Select the name of the volume that you want to restore and click **Restore**. A dialog box opens, asking you to confirm that you want to restore the volume.
5. Make sure that you do not have another volume of the same name in your list of volumes. If you have one, change the name of the volume you want to restore and then click **Yes**.

   - To restore the volume without changing the name, click **Yes**.
   - If the volume name provided does not exist, the system will restore the volume with the default name created in the recovery bin (in the CLI) and will not prompt you for another name.

   NOTE: Recovered volumes are in an offline state after they have been recovered. You must set the volume online before you can use it.

Some volume information is not restored when a deleted volume is recovered, including:

- Membership in a volume collection
- RAID preferences
- Scheduling information
- Snapshot reserve (set to 0 when the volume is deleted)
- Snapshots for that volume
- SyncRep status
- Replicated volumes

You can reconfigure this information for a restored volume.

Purge Deleted Volumes

Purging a deleted volume permanently removes a volume and all its data from the recovery bin. Purged volumes can never be recovered.

When a volume is deleted and appears in the recovery bin, it remains there for up to 1 week (depending on available free space or on group limits, such as volume count) before the firmware automatically purges it. However, you can also manually purge volumes from the recovery bin.

1. Click **Volumes**.

2. Click **Volumes** in the Volumes panel (not an individual volume name).
3. In the Activities panel, click **Manage recovery bin**. The Volume Recovery Bin dialog box opens.
4. Select the volume name that you want to permanently delete in the recovery bin.
5. Click **Purge**.

   To purge all volumes in the recovery bin at the same time, click **Purge All**.
6. When prompted to confirm the decision, click **Yes** to continue with the purge or **No** to cancel.

NOTE: When a volume is purged, all of its data is lost.

About Changing the Reported Volume Size

You can increase the reported size of a volume while the volume is online and without disrupting access to the volume. To decrease the size of a volume, it must first be set offline.

Changing the size of a volume is subject to the following conditions:

- Before changing a reported volume size, Dell recommends that you fully understand the impact on the operating system, file system, and applications using the volume.
- Not all operating systems, file systems, and applications easily handle volume size changes or behave in a predictable manner when you change a volume size.
- You cannot change the size of a template volume.

Changing the reported volume size affects the space that the group allocates to the volume and for volume snapshots and replication as follows:

- For a volume that is not thin-provisioned, changing the reported size proportionally changes the amount of space the group allocates to the volume (the volume reserve).
- For a thin-provisioned volume, changing the reported size changes the minimum volume reserve, in-use space warning limit, and maximum in-use space, because they are based on a percentage of the reported size. The space that the group allocates to the volume (the volume reserve) might also change.
- If the volume reserve changes due to the reported volume size change, snapshot space and replication space also change.

NOTE: If you are replicating the volume, the secondary group does not recognize the reported size change until the next replication.

Decrease the Reported Size of a Volume

Decreasing the size of a volume is sometimes called “shrinking” a volume.

The following considerations apply to decreasing the size of a volume:

- The volume must be set offline before its size can be decreased. After you finish decreasing the volume size, you can set the volume online again.
- ⚠ **CAUTION: If you decrease a volume size to less than the amount of space currently in use, you could lose data.**
- If you do not specify a unit for the size, the unit defaults to MB.
- Resizing out-of-sync synchronous replication (SyncRep) volumes is not supported.

You cannot use the Group Manager GUI to decrease the reported size of a volume. Instead, you must use the following Group Manager CLI command:

```
volume select volume_name shrink new_size
```

See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about using CLI commands that relate to volumes.

Increase the Reported Size of a Volume

You can increase the reported size of the volume while the volume remains online.

The following considerations apply to increasing the size of a volume:

- If the size you specify is not a multiple of 15MB, the group rounds up the value to the nearest multiple of 15MB.
- If you do not specify a unit for the size, the unit defaults to MB.
- If you configured the volume for replication, the wizard shows the delegated space on the replication partner.
- If the new volume size exceeds the capacity of a pool, the free space cell displays a negative value.
- For thin-provisioned volumes, you can modify the in-use warning value and maximum in-use space value.
- Resizing out-of-sync synchronous replication (SyncRep) volumes is not supported.

To increase the reported volume size:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **Space** tab to open the Modify Volume Settings – Space dialog box.
5. Specify the new reported volume size in the **Volume size** field.

   (Values in the Pool Space table change, based on the new volume size.)
6. (Optional) Use the slider bars to modify the in-use warning value and maximum in-use space value for a thin-provisioned volume.
7. Click **OK**.
8. Confirm that you want to create a snapshot of the current volume before you resize it.

You can also increase the size of a volume with the Group Manager CLI volume select size command. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about using CLI commands that relate to volumes.

About Reclaiming Unallocated Space

Version 6.0 of the PS Series firmware introduced support for using SCSI unmap operations to recover unused space previously allocated to volumes.

As a host writes data to a volume, the array allocates additional space for that data. Prior to support of the volume unmapping feature, that space remained allocated to the volume, even if the data was deleted from the volume and hosts were no longer using it. This usage created a “watermark” effect; the array could not deallocate the space, and therefore the space was unavailable for allocation to other volumes.

With volume unmapping, the array can reclaim this unused space. When a host connected to a volume issues a SCSI unmap operation, the array deletes the data and deallocates the space, making it available for allocation by other volumes.

Volume unmapping requires that all group members be running version 6.0 or later of the PS Series firmware. Space deallocation occurs only if you are using initiators that are capable of sending SCSI unmap operations and only on a “best effort” basis. Although space deallocation occurs on both thin-provisioned volumes and regular volumes, the volume reserve size can potentially shrink as a result of unmap operations only on thin-provisioned volumes.

If space is deallocated, that space might not be immediately available as free space for other volumes until the array deletes snapshots of those volumes.

Unmap operations are supported in the following operating systems:

- VMware ESX 5.0
- Red Hat Enterprise Linux 6.0

- Windows 8/Windows Server 2012

Running Defrag Tools

Run defragmentation tools (such as fstrim, windows manual defrag, or esxtool) during periods of low I/O activity, because these operations might result in large numbers of unmapping operations and reduce array performance.

Unmapping Replicated Volumes

You should not run operations that result in SCSI unmap operations (for example, format or defrag) on volumes on which replication (including synchronous replication) is enabled.

Such operations on replicated volumes result in zeros being written to the destination sectors of the volume. This writing of zeros can cause the operations to take a long time to complete, and no space is reclaimed. Therefore, Dell recommends that you disable unmapping on hosts or volumes (depending on the host operating system) that are using replicated volumes. Refer to the following sections for information about disabling and enabling unmapping in VMware ESX, Red Hat Enterprise Linux, or Windows 8/Windows Server 2012 operating systems.

Using Unmapping with VMware

In VMWare ESX 5.0, automatic unmapping is enabled by default for deletion operations such as deleting or migrating a VM.

However, Patch ESXi500-201112001, a recommended patch for ESX 5.0, disables automatic unmapping. If you are running this patch, you can enable unmapping using the `-y` argument to the `vmkfstools` utility, which performs manually invoked unmapping and space reclamation on a per-volume basis.

For more information about using `vmkfstools`, refer to the VMware Knowledge Base.

Using Unmapping with Red Hat Enterprise Linux

By default, unmapping is disabled in Red Hat Enterprise Linux.

Linux file systems that support unmap operations, such as ext4 on Red Hat Enterprise Linux 6, must be mounted with the`-o discard` option to enable free space recovery functionality. If the file system does not support the`-o discard` mount option, it does not support free space recovery on PS Series storage array volumes.

When free space recovery using volume unmapping is enabled, larger files are deleted and their space is reclaimed. However, volume unmapping operations can stop new writes to the volume and adversely affect performance. This scenario is most likely with replicated volumes. Therefore, Dell recommends that you not use volume unmapping with replicated volumes. Also, running `mkfs`, defragging, or deleting files on file systems mounted using the `-o discard` option can generate large numbers of unmap operations, which are not recommended for replicated volumes.

Using Unmapping with Windows 8/Windows Server 2012

By default, unmapping is enabled in Windows 8/Windows Server 2012, which automatically detects a volume’s provisioning capabilities, including whether or not the volume can process unmap operations. When data is deleted from a volume in Windows, the corresponding space on the PS Series storage array volume is also freed automatically, thus maximizing the efficiency of the array.

To disable unmapping in Windows 8/Windows Server 2012, issue the following command:

```
fsutil behavior set disabledeletenotify 1
```

To reenable unmapping, issue the following command:

```
fsutil behavior set disabledeletenotify 0
```

To check the current setting of unmapping, issue the following command:

```
fsuitl behavior query disabledeletenotify
```

NOTE: The `disabledeletenotify` setting is a global operating system setting that not only disables unmap operations from being sent to the PS Series storage arrays, but also disables TRIM to SSDs. For more information about the `fsutil` utility, see: technet.microsoft.com/en-us/library/cc785435(v=ws.10).aspx

Automatic defrag can generate a large number of unmap operations, which are not recommended for replicated volumes. Therefore, in Windows Server 12, turn off automatic defrag for replicated volumes.

Set a Volume or Snapshot with Lost Blocks Online

In rare circumstances, a volume (or snapshot) might lose blocks. For example, loss can occur if the power fails and then a control module cache battery fails. (If the control module cache battery is the only power source for a control module for more than 72 hours after a power failure occurs, the battery can fail.)

If a volume (or snapshot) loses blocks, the current status of the volume (or snapshot) is `offline-lost-cached-blocks`. In addition, the group generates an event message.

You can choose to set the volume online but retain the lost blocks. If an application tries to read a lost block, an error occurs. If an initiator writes new data to a lost block before it is read, the block is no longer lost. The members containing lost blocks have a status of `RAID lost blocks` until initiators write to all the lost blocks.

NOTE: Setting a volume with lost blocks online is a data integrity risk. The blocks might contain old or invalid data.

To manage a volume or snapshot with lost blocks:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click the **Status** tab.
4. Click **offline-lost-cached-blocks**.
5. Either:
   - Click **Set the volume online but retain the lost blocks** to set the volume or snapshot online but keep the lost blocks. The volume (or snapshot) status changes to `online-lost-blocks`.
   - Click **Mark the lost blocks valid and set the volume online** to set the lost block status to valid. The volume (or snapshot) status changes to `online`. The status of the members containing volume data changes to `online`.
6. Click **OK**.

Volume and Snapshot Status

Table 35. Current Volume and Snapshot Status shows status values for volumes and snapshots, and provides solutions for any problems.

Table 35. Current Volume and Snapshot Status

Status	Description	Solution
`online`	Administrator set the volume or snapshot online, and no failures have occurred.	None needed; informational. Online volumes and snapshots are shown in the far-left panel in black text.
`offline`	Administrator set the volume or snapshot offline.	None needed; informational. Computers cannot access the volume or snapshot, but no failures have occurred. Offline volumes and snapshots are shown in the far-left panel in gray text.

Status	Description	Solution
`offline (snap reserve met)`	Volume or snapshot was automatically set offline due to the selected snapshot recovery policy.	Increase the amount of reserved snapshot space.
`offline (thin max grow met)`	A thin-provisioned volume and its snapshots were automatically set offline because a write exceeded the maximum in-use space value.	Increase the value of the maximum in-use space setting or increase the volume's reported size.
`offline (missing pages)`	A volume or snapshot was set offline because some volume data cannot be found. This condition is serious.	Contact your PS Series support provider.
`offline (nospace auto grow)`	The thin-provisioned volume and its snapshots were set offline because not enough free pool space existed for the volume reserve to increase automatically.	Increase pool free space. For example, you can add another member to the pool or move volumes from the pool.
`offline (member down)`	Volume or snapshot was automatically set offline because a member that contains volume or snapshot data is unavailable.	Identify why the member is unavailable and correct the problem.
`offline (lost blocks)`	Volume or snapshot was automatically set offline because blocks were lost. Computers cannot access the volume or snapshot. This condition is serious.	Click the status link and select how to manage the lost blocks.
`unavailable due to SyncRep`	Volume or snapshot is unavailable because SyncRep is confused, often due to a mismatch between the state in the database and the expected state.	Determine the reason for the issue and correct the problem.
`available (no new connections)`	During a failover, SyncRep is trying to determine which volume should be active. Until then, no new connections are allowed.	Wait for SyncRep for to determine which volume is active.
`unavailable due to internal error`	Volume or snapshot is unavailable due to an unexpected error.	Contact Customer Support for assistance.

Volume and Snapshot Requested Status

Table 36. Volume and Snapshot Requested Status shows the possible values for the requested status for a volume or snapshot.

Table 36. Volume and Snapshot Requested Status

Status	Description
`online`	Administrator set the volume or snapshot online.
`offline`	Administrator set the volume or snapshot offline. Computers cannot access an offline volume or snapshot.
`online (lost blocks)`	Administrator set the volume or snapshot online despite lost blocks. Authorized computers can access the volume or snapshot. If an application tries to read a lost block, an error occurs. If the block is rewritten, no error occurs, and the block no longer shows a status of lost.

About Managing Storage Capacity Utilization On Demand (Thin Provisioning)

You can use thin-provisioning technology to more efficiently allocate storage space, while still meeting application and user storage needs. With a thin-provisioned volume, the group allocates space based on volume usage, enabling you to “over-provision” group storage space (provision more space than what is physically available).

However, if your environment requires guaranteed space for volume, thin provisioning might be inappropriate. Thin provisioning is most effective if you can accurately predict how volume usage increases over time.

NOTE: Dell recommends that you fully understand the benefits and risks of using thin provisioning before implementing it in your environment. Environments that use thin provisioning should have around-the-clock support to handle any space allocation issues and prevent service-level disruption.

Thin provisioning volumes is beneficial in a number of environments. For example, if your environment does not easily allow you to expand file systems or raw disks, you can give thin-provisioned volumes excessively large reported sizes to account for future growth. The group automatically allocates space to volumes only if usage patterns warrant the space.

Thin provisioning also helps you plan for future group expansion. For example, you can size volumes according to their maximum possible space requirements, even if the group currently cannot provide all the required space. As volume usage increases, you can expand group capacity, with no user impact. You do not need to change drive letters, expand volume sizes, or add volumes.

When you create a volume, you specify the reported size for the volume. The reported size is seen by iSCSI initiators. The actual amount of pool space that the group allocates to a volume is called the volume reserve. The value of the volume reserve depends on whether you enable thin provisioning on a volume:

- Thin provisioning disabled

  The volume reserve is equal to the reported size.

  For example, even if only 10 percent of a volume is in use, the group allocates the full reported size.
- Thin provisioning enabled

  The volume reserve is equal to or less than the reported size, depending on volume usage and the thin-provisioning settings.

Initially, the group allocates the minimum amount of volume reserve for a thin-provisioned volume. The minimum is 10 percent of the reported volume size or the user-specified percentage.

As initiators write to the volume, free volume reserve decreases. When free volume reserve falls below a threshold, the group increases volume reserve, up to a user-defined maximum (assuming available free pool space):

- For a volume with a reported size of 100GB or greater, when free volume reserve is less than 6GB, the group allocates an additional 10GB.
- For a volume with a reported size that is less than 100GB, when free volume reserve falls below 6 percent of the reported volume size, the group allocates an additional 10 percent of the reported volume size.

Event messages inform you when in-use volume reserve surpasses a user-defined limit and reaches the maximum.

Enable Thin Provisioning on a Volume

When you create a new volume or clone an existing volume, you can enable thin provisioning on the volume. In addition, you can modify an existing volume and enable thin provisioning. You can make these changes while the system is running.

Thin provisioning is not appropriate for all environments or volumes. You must fully understand thin provisioning before implementing the functionality on a volume.

When enabling thin provisioning on an existing volume, the following considerations apply:

- Enabling thin provisioning on a volume usually decreases the amount of space that the group allocates to the volume (called the volume reserve).
- Enabling thin provisioning changes the amount of allocated snapshot space and replication space, because the group allocates snapshot space and replication space based on a percentage of the volume reserve. However, the group increases the snapshot space and replication space percentages to prevent the deletion of snapshot or replication data.

To enable thin provisioning on an existing volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **Space** tab to open the Modify Volume Settings – Space dialog box.
5. Select **Thin provisioned volume**.

   The Pool Space table values change.
6. (Optional) Modify the groupwide default thin-provisioning space settings:
   - Minimum volume reserve
   - In-use space warning limit
   - Maximum in-use space (maximum volume reserve)
7. Click **OK**.

Make sure that you carefully monitor the space usage for a thin-provisioned volume.

Thin-Provisioning Space Settings

Three settings control how the group allocates space to thin-provisioned volumes and when the group generates events related to space usage:

- Minimum volume reserve

  Minimum amount of pool space that the group allocates to the volume, based on a percentage of the reported volume size. The default groupwide setting is 10 percent.
- In-use space warning limit

  Amount of in-use volume reserve that results in notification, based on a percentage of the reported volume size. The default groupwide setting is 60 percent.

  When in-use volume reserve reaches the in-use warning limit, the group generates a warning event message. Additional warning event messages occur as follows:

  – For volumes larger than 200GB, when the in-use volume reserve increases by every additional 10GB.

  – For volumes smaller than 200GB, when the in-use volume reserve increases by every additional 5 percent.

  For example, if you create a thin-provisioned volume with a size of 500GB and set the warning limit to 75 percent, a warning occurs when the amount of in-use volume reserve is more than or equal to 75 percent of 500GB, or 375GB.
- Maximum in-use space

  Maximum amount of in-use volume reserve (maximum size of the volume reserve), based on a percentage of the reported volume size. The default groupwide setting is 100 percent.

  The maximum in-use space value determines the behavior when the volume reserve reaches its maximum size:

  – If the maximum in-use space value is less than 100 percent, and an initiator write exceeds this limit, the write fails. The group sets the volume offline and generates event messages.

  If you increase the maximum in-use space value or the reported volume size (both operations require free pool space), the group automatically sets the volume online and writes succeed.

- If the maximum in-use space value is 100 percent, and an initiator write exceeds this limit, the volume is not set offline; however, the write fails, and the group generates event messages. If you increase the reported size of the volume, writes succeed.

  This behavior is the same as when in-use space for a volume that is not thin-provisioned reaches its reported size.

The maximum in-use space value helps prevent one volume from consuming all the pool free space and setting other thin-provisioned volumes offline.

You can change the groupwide default volume settings and override the default values when you create a thin-provisioned volume, or modify a volume and change the settings.

Modify the Thin-Provisioning Space Settings

To modify the thin-provisioning space settings:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Modify settings**.
4. Click the **Space** tab to display the sliders for changing the settings.

   Make sure the **Thin provisioned volume** checkbox is selected. Otherwise, you will not be able to see the sliders.
5. Move the sliders to adjust (in this order) the settings for:

   - Minimum volume reserve
   - In-use space warning limit
   - Maximum in-use space (maximum volume reserve)

   The values in the table change based on the new values. If a change exceeds capacity, the **Free pool space** field contains a negative value.
6. Click **OK**.

Disable Thin Provisioning on a Volume

You can disable thin provisioning on a standard volume.

NOTE: You cannot disable thin provisioning on a template volume, thin clone volume, recovery template volume, or a recovery thin clone volume.

When disabling thin provisioning, the following considerations apply:

- If you disable thin provisioning on a volume, the group allocates the full reported volume size (the reported size and the volume reserve is the same). Therefore, you must have sufficient free pool space.
- Because the group bases snapshot space and replication space on a percentage of the volume reserve, disabling thin provisioning increases snapshot space and replication space. Therefore, you must have sufficient free pool space.
- In some cases, if you disable thin provisioning on a volume, the group automatically decreases the snapshot reserve percentage to prevent an excessive allocation of snapshot space.

  Excessive allocation can occur if you previously set the snapshot reserve percentage to a high value to prevent the group from deleting snapshots (for example, if you increased the snapshot reserve percentage from 100 percent to 500 percent). If you disable thin provisioning on the volume, the group might decrease the percentage to a more appropriate value, closer to 100 percent. The group does not decrease the snapshot reserve percentage if the decrease requires deleting snapshots.

To disable thin provisioning on a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Modify settings**.
4. Click the **Space** tab to open the Modify Volume Settings – Space dialog box.
5. Deselect **Thin-provisioned volume**.

The Pool Space table values change, based on the new volume setting. If the volume change exceeds pool capacity, the free space field displays a negative value.

6. Click **OK**.

About Improving Pool Space Utilization (Template Volumes and Thin Clones)

Some computing environments use multiple volumes that contain a large amount of common data. For example, some environments clone a standard volume and create multiple “boot volumes” that administrators use to boot different client computers. Most of the data is common to all the volumes; only a small portion of volume space contains unique data. Because each boot volume consumes pool space for the common data, the group is storing multiple copies of the same data, which is not an efficient utilization of space.

To use pool space more efficiently, instead of cloning standard volumes, you can create one volume and populate it with the common data. After you convert the volume to a template volume, you can create multiple thin clone volumes and then write to each thin clone to make it unique. For example, you can add data such as a page file to a thin clone.

Because a thin clone shares the common data in the template volume, each thin clone only consumes the space needed to store the differences (or deltas) between the thin clone and the template volume.

Initially, a template volume and thin clone are identical in reported size and content. Because the group allocates space to the new thin clone in the same way it allocates space to a new standard, thin-provisioned volume, only the minimum volume reserve is consumed from free pool space.

When initiators write data to a thin clone, space is consumed from free volume reserve. As needed, the group allocates additional volume reserve to the thin clone, up to the maximum in-use space setting for the thin clone.

You can also modify the thin clone and change the data that the thin clone shares with the template volume. However, the data in the template volume is always preserved because a template volume is read-only. Group Manager tracks the amount of data that is shared between each thin clone and template volume.

With a few exceptions, all normal volume operations apply to template volumes and thin clones.

Keep in mind:

- Thin clones are considered attached to the template volume and cannot exist without the template volume, similar to how snapshots depend on the base volume.
- The group always maintains and shows the dependency of a thin clone on a template volume. You can use the GUI to display all volumes in alphabetical order or display thin clones under the template volume on which they depend.
- When you replicate a template volume and its attached thin clones, the primary and secondary groups maintain the dependency. For example, you must replicate a template volume before replicating any of its thin clones. On the secondary group, you can display thin clone replica sets under the template replica set on which they depend.
- The group maintains and shows the dependency of a thin clone on a template volume (or a thin clone replica set on a template replica set), even if a volume (or replica set) changes state, as occurs during failover and failback operations. For example, if you promote a thin clone replica set to a recovery thin clone, you can still see the dependency of the recovery thin clone on the template replica set.
- You cannot delete a template volume, a template replica set, or a recovery template if a thin clone, thin clone replica set, or recovery thin clone depends on it.
- You cannot disable replication on a template volume until you disable replication on all its thin clones.

Space Considerations for Template Volumes and Thin Clones

When you convert a standard volume to a template volume:

- Thin provisioning is enabled on the volume, the volume is set offline, and the volume permission is set to read-only.

NOTE: If you are using the Group Manager CLI, you must perform these tasks manually before you can convert to a template volume.

- Volume reserve decreases to the amount of in-use space (or the minimum volume reserve, whichever is greater), and free volume reserve becomes unreserved space.
- Snapshot reserve is adjusted, based on the new volume reserve. If necessary to preserve existing snapshots, the snapshot reserve percentage is increased.

When you create a thin clone volume, it has the same reported size and contents as the template volume. If you mount the thin clone, you can see the data that the thin clone shares with the template volume.

The group allocates only the minimum volume reserve when you first create a thin clone. The group allocates additional space if you specify snapshot reserve for the thin clone. As with a standard, thin-provisioned volume, as you write to a thin clone, the group allocates more space and increases the volume reserve as needed.

In the Volume Status window for a thin clone volume, the Volume Space table shows the space utilization for the thin clone, including the in-use space, which is the portion of volume reserve that is storing data unique to the thin clone. When you first create a thin clone, it has zero in-use space.

Also in the Volume Status window for a thin clone volume, the Shared Space table (appears only for thin clones) shows the amount of space that is shared with the template volume and the unshared (in-use) space. As you write to the thin clone, unshared (in-use) space increases. In some cases, when you write to a thin clone, shared space can decrease (for example, if you are overwriting shared data).

Free space in the Shared Space table shows the amount of unwritten thin clone space (that is, the reported volume size minus the combined shared space and unshared space). This space represents the amount of data you can write to the thin clone before you need to increase its size. This value is the same as the value for “unreserved” space in the Volume Space table in the Volume Status window for the template volume.

If you detach a thin clone, the resulting new standard volume has in-use space equal to the combined shared space and unshared space, as shown in the Shared Space table in the Volume Status window.

Restrictions on Template Volumes and Thin Clones

With a few exceptions, all normal volume attributes and operations apply to template volumes and thin clones as specified in Table 37. Template Volume and Thin Clone Restrictions.

Table 37. Template Volume and Thin Clone Restrictions

Attribute or Operation	Restriction
Snapshots	You cannot restore a template volume from a snapshot.
Thin provisioning	You cannot disable thin provisioning on a template volume or thin clone.
Permissions	You cannot set template volumes to read-write permission.
Volume collections	You cannot include a template volume in a volume collection.
Scheduling operations	You cannot schedule a snapshot or replication operation for a template volume.
RAID preference	Thin clones inherit the RAID preference, if any, of the template volume.
Member binding	Thin clones inherit the member binding setting, if any, of the template volume.
Cloning	Cloning a template volume creates a new template volume with a new name and iSCSI target, but the same reported size, pool, and contents as the original volume at the time of the cloning. Cloning a thin clone creates a new thin clone with a new name and iSCSI target, but the same reported size, pool, contents, and relationship to the template volume as the original thin clone at the time of the cloning.
Resizing	You cannot change the reported size of a template volume. However, you can change the thin-provisioning settings.

Attribute or Operation	Restriction
Pool move	Thin clones inherit the pool setting of the template volume. If you move the template volume to a different pool, the thin clones also move.
Replication	You can replicate a template volume only one time. You cannot replicate a thin clone until you replicate the template volume to which the thin clone is attached.
Failover	You can permanently promote a template replica set to a template volume only if you first permanently promote all the attached thin clone replica sets to thin clones.
Failback	You cannot demote a template volume to a failback replica set. You cannot fail back a template volume. To fail back a thin clone volume, the template volume must exist on the primary group.
Deletion	You cannot delete a template volume if it has thin clones or failback thin clone replica sets attached to it. You cannot delete a recovery template volume if any recovery thin clone volumes, thin clone replica sets, or permanently promoted thin clone replica sets are still attached to the volume.
Synchronous replication	See About Using Thin Clones and Templates with Synchronous Replication.

Convert a Standard Volume to a Template Volume

You can convert standard volumes to a template volume. The following prerequisites and considerations apply:

- When you convert a standard volume to a template volume, the template volume is thin provisioned, read-only, and offline. You can set the volume online at any time.
- When you convert to a template volume, the group disables any schedules that include the volume. If you later convert the template volume to a standard volume, the group does not automatically enable the schedules.
- Before converting to a template volume, make sure that the standard volume contains all the data that is shared with the thin clones.
- You cannot convert a volume to a template while it is a member of a collection.
- Make sure that the standard volume has sufficient free space to hold the approximate amount of data that you write to each thin clone.

  For example, if the reported size of the template volume is 1GB, and in-use space is 900MB, you can write approximately 100MB to each thin clone before you must increase the thin clone size.

To convert a standard volume to a template volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Convert to template**.
4. Confirm that you want to convert the volume to a template volume.

Create a Thin Clone

To create thin clones from template volumes, you need the following information:

- Clone name

  Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- Clone description (optional)

  Descriptions can be up to 127 bytes long. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- Storage space assignment:

  - Snapshot reserve setting
  - Thin-provisioning setting
- Access control
- Multihost access

To create a thin clone from a template volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the template volume.
3. Click **Create thin clone** to open the Create Thin Clone – Volume Settings dialog box.
4. Type a unique name and, optionally, a description.
5. Click **Next** to open the Create Thin Clone – Space dialog box.
6. Change the snapshot reserve setting and the thin-provisioning settings.

   NOTE: If data reduction has been enabled on the volume, snapshot reserve is permanently disabled.
7. Click **Next** to open the Create Thin Clone – iSCSI Access dialog box.
8. Create an access control policy for the volume, select the permission (read-write or read-only), and specify the multihost access setting.
9. Click **Next** to open the Summary dialog box.
10. Click **Finish** if the thin clone configuration is correct. Click **Back** to make changes.

Detach a Thin Clone from a Template Volume

You can detach a thin clone from its template volume. Detaching a thin clone from a template volume breaks the dependency between the thin clone and the template volume.

The following restrictions apply:

- You cannot detach a thin clone if replication is enabled for the thin clone.
- If you detach a thin clone from a template volume, the thin clone is converted to a standard volume and no longer shares space with the template volume.
- The volume reserve for the thin clone increases by the amount of space the thin clone shares with the template volume.
- When you detach a thin clone from a template volume that is bound to a member or has a RAID preference, the resulting volume does not inherit the binding or the RAID preference.

To detach a thin clone from a template volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the thin clone.
3. Click **Detach from template**.
4. Confirm that you want to detach the thin clone.

Convert a Template Volume to a Standard Volume

You can create template volumes and then later convert them to standard volumes, with the following restrictions:

- You cannot convert a template volume to a standard volume if thin clones are attached to the template volume.
- If a template volume is configured for replication, you must disable replication before you convert a template volume to a volume.
- Space used to store template replicas on the secondary group becomes unmanaged if you convert a template volume to a standard volume.

To convert a template volume to a standard volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.

3. Click **Convert to volume**.
4. Confirm that you want to convert the template volume to a standard volume.

About Data Center Bridging

Data center bridging (DCB) is a set of extensions to IEEE Ethernet standards, intended to improve the performance, controllability, and reliability of the Ethernet fabric. DCB can make Ethernet the unified fabric for different types and priorities of data traffic in the data center.

The DCB enhancements are aimed to:

- Improve Ethernet performance
- Reduce or eliminate lost packets for lossless traffic classes
- Provide the ability to define classes of traffic, which can have differing bandwidth, priority, and frame loss characteristics, enabling:
  - Pause per class — Ability to independently pause based on user priorities or classes of service
  - Bandwidth allocation or traffic class prioritization — Ability to provide different service characteristics to traffic classes
  - End-to-end congestion management — Ability to notify sender of rate limiting transmission to avoid frame loss
- Enable deterministic performance during congestion
- Provide a unified fabric for converged SAN and LAN for improved scalability and performance.

NOTE:

- All group members must be running PS Series firmware version 5.1 or later to use DCB.
- All group members must be 10Gb-based arrays. Mixed groups of 1Gb and 10Gb arrays are not supported.

Configure Data Center Bridging

Ensure that you have properly configured your environment (switch fabric) for all DCB-enabled devices, according to the requirements described in the following section.

Requirements

For DCB to function correctly, your configuration must meet the following requirements:

- All group members must be 10Gb arrays.
- All group members must be running PS Series firmware version 5.1 or later.
- Initiator and target devices, as well as all the network switches between the initiators and the targets, must support DCB. All devices must be configured for DCB from end to end, including:
  - From the host's CNA card (edge-port) to the network switch
  - From the network switch to the array (edge-port)
  - Everything else in between (if using multiple switches)
- All of the network devices, initiators, targets (if non-EQL targets), and switches in the network configuration, such as PFC (Priority Flow Control), iSCSI TLV (Type-Length-Value), and ETS (Enhanced Transmission Selection), must be supported.
- Even if all the attached network switches fully support DCB, you must configure DCB properly on the network switches for the configuration to work correctly.

Additional Information

Dell strongly recommends that you read the following white papers for more information about configuring DCB:

- The *Data Center Bridging: Standards, Behavioral Requirements, and Configuration Guidelines with Dell EqualLogic iSCSI SANs* white paper contains information about the DCB requirements and configuration guidelines. You can download the paper from: http://en.community.dell.com/techcenter/storage/w/wiki/4396.data-center-bridging-standards-behavioral-requirements-and-configuration-guidelines-by-sis.aspx
- The *EqualLogic DCB Configuration Best Practices* white paper contains the DCB requirements for EqualLogic, to ensure that DCB is properly enabled or disabled across all devices, and to assist with identifying and resolving basic DCB configuration issues

in the EqualLogic SAN. You can download the paper from: http://en.community.dell.com/techcenter/storage/w/wiki/4355.configuring-dcb-with-equallogic-sans.aspx

Set the Data Center Bridging VLAN ID

To set the DCB VLAN ID for the first time:

1. Click **Group** → **Group Configuration**.
2. Click the **Advanced** tab.
3. Enter the DCB VLAN ID in the **VLAN ID** field; it must be a value from 0 to 4095. This value must be the same as the VLAN ID configured in the switch for use by the iSCSI SAN. Values 0 and 1 are not recommended.

NOTE:

- **The VLAN IDs on the switch must match the VLAN you specify on the group. The ports on the switch need to allow both the old VLAN (which usually implies a VLAN ID of 0 if none has been set) and the new VLAN ID or all VLANs (trunk mode) to enable the group to communicate with the switch.**
- **All group members must be 10Gb arrays and be active.**
- **You can also specify the DCB VLAN ID through the CLI using the grpparams dcb def-vlan-id command. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.**

Change the Data Center Bridging VLAN ID

If your group's DCB VLAN ID has been changed from the default value of 2, you must configure the group to use it.

NOTE:

- All group members must be active.
- This procedure requires modifying the configuration of third-party products. See the documentation for your switches for detailed instructions on these tasks.

1. Set the switch ports used by group members to allow both the old and new VLAN IDs.
2. Bridge the old and new VLANs either by modifying switch settings or by creating an untagged access port for each VLAN and bridging the connections using the appropriate cable.
3. Set the VLAN ID in the group:
   a. Click **Group** → **Group Configuration**.
   b. Click the **Advanced** tab.
   c. In the Network Management panel, enter the new VLAN ID in the VLAN ID field. The value must be from 0 to 4095.
4. Reconfigure iSCSI initiators to use the new VLAN ID.
5. Unbridge the old and new VLANs.
6. (Optional) Modify switch ports so that the old VLAN ID is no longer allowed.

Monitor Data Center Bridging Statistics and Configuration

1. Click **Group**.
2. Expand **Members** and then select the group member.
3. Click the **Network** tab.
4. In the IP Configuration panel, select a network interface.
5. The Switch Information section indicates the LLDP state of the switch connected to the interface.
6. In the Activities panel, click **DCB Details** to open the dialog box.

NOTE: The DCB Details link appears dimmed if the member does not support DCB.

10

VMware Group Access Panel

The VMware Group Access panel allows you to view VMware virtual volume (VVol) settings, if VVols are configured for your group.

To open the VMware Group Access panel, click the **VMware** tab in the selection pane on the lower-left side of the Group Manager window.

NOTE: If the VMware tab is not displayed in the selection pane, move the View Drag handle (the two arrows above the list of tabs) up to expand the list. Or, click the VMware icon (🖥).

When the VMware Group Access panel is first opened, it shows the **Overview** page.

VMware Overview Panel

When you click the VMware tab (or VMware icon) in the selection pane, the **Overview** page opens.

The Overview page contains the following information:

- Storage Containers
  - Total count of storage containers
  - Total physical size allocated to storage containers
- VVols
  - Total number of VMs
  - Total number of VVols
  - Total number of VVol snapshots
  - Total number of VVol linked clones

The navigation tree provides access to the following items:

- Configuration – Opens the **Protocol Endpoints Access Control List** page
- Storage Containers – Opens the **Storage Containers** page and shows details for existing storage containers
- Virtual machines – Opens the **Virtual Machines** page, which shows all virtual machines that are resident on an array, the VVols that each is composed of, and the storage containers that the VVols exist in

The lifecycle of the virtual machine (VM) is controlled using VMware vSphere and on EqualLogic environments are managed by the Dell Virtual Storage Manager for VMware (VSM). VSM creates the VM entities on the array, using actions triggered by vCenter and ESXi either directly or without user invocation.

NOTE: Although you can use the Group Manager to create storage containers, Dell recommends that you use the Dell Virtual Storage Manager (VSM) to perform this task. The VSM procedures for setting up VVols is simpler than using the Group Manager. Using the Group Manager for VVol operations is recommended for experienced users only.

For detailed information about setting up VVols using VSM, see the *Dell Virtual Storage Manager for VMware Installation and User's Guide*.

About Protocol Endpoints

A protocol endpoint is the iSCSI target used for VVol storage, and the mean by which to access VVol storage containers. In order to perform VVol user operations from within vCenter, protocol endpoint access rules need to be established.

NOTE: Configuring a protocol endpoint requires establishing one or more access policies. Dell requires the use of Virtual Storage Manager (VSM) to establish Protocol Endpoint (PE) access rules.

To view all access policies defined for a group:

1. Click the **Group** tab.
2. Click **Group Configuration** → **Acess Policies**.

   Group Manager displays a table of all access policies.

By selecting an access policy in this table, you can view the associations for that policy, whether that association is ”protocol endpoint” or a “legacy.”

See About Access Policies.

About Storage Containers

A storage container is a quantity of storage made available for the placement of virtual volumes (VVols). The VVol architecture defines a storage subsystem consisting of a VASA Provider (VP) and some storage devices. The VP acts as the control plane intermediary between the VMware ESX virtual machine server platform and/or the VMware vCenter management platform and the storage devices, which equates to one or more PS Array groups.

NOTE: Storage containers cannot be used to hold standard EqualLogic volumes, snapshots, or NAS containers. Storage containers are not supported outside of the virtual volumes context.

In the virtual volume architecture, all physical and logical storage in the storage container is treated equally. Within a storage container, all types of VVols consume storage from a single, undivided free pool. VVols and their derivatives (snapshots, clones) all draw from the container’s space in equal terms, and their full logical size is debited from the container’s logical free space when created.

While users can manage the amount of physical space, no restrictions apply as to how the physical space is used within the container. Snapshots, clones, and VVols share the space on a first-come, first-served basis.

Storage Container Limitations

- No specific limit applies to the number of VVols in a single storage container; however, the number is limited by the space the VVols consume and the physical and logical storage size of the storage container.
- A storage container’s physical size is limited only by the size of its parent pool.
- Storage container names must be between 2 and 42 alphanumeric Unicode characters. They must begin with a letter, and can contain semicolon, period, and hyphen characters.

  NOTE: While storage container names are case sensitive in terms of display format, two containers cannot have names that differ only by case. For example, if a storage container name “VvolSC” exists, new containers named “vvolsc” or “vvolSc” cannot be created.
- A storage container is located on each member of its parent pool.
- Deleting a pool with an existing storage container is supported. Merging two pools with existing storage containers is supported. Moving a storage container to a different pool is not supported.
- Storage containers cannot be explicitly moved between pools. If two pools are merged via either the “pool merge” or “pool delete” operation, any Storage containers that draw from the deleted or merged pool will be automatically reassigned to the default or the merged-to pool respectively.

Storage Container Space Limits

Storage container space is managed by the storage administrator. Storage administrators must allow access to enough physical storage to meet the needs of a storage container of a given logical size under a reasonable range of operating circumstances, without preallocating too much physical storage to the container.

A storage container will have access to a finite amount of physical storage, which is defined as the container's physical size.

It is possible that a storage container could run out of logical space long before it runs out of physical space. This could occur if there is an aggressive snapshot schedule which quickly consumes logical space. Logical space is a function of the size of all VVols existing, independent of whether it is written. Resizing the storage container (datastore) to be larger using the VSM user interface allows for an increase of both physical and logical space. In the VSM user interface, you can monitor both the logical and physical space consumption of the datastore/storage container by selecting the specific datastore, then **Manage → Dell VSM**.

Storage container accounting is done in units of logical space. That is, the storage container size is specified in terms of MB, GB, or TB. The logical space that each VVol within the container consumes is always equal to the nominal size of the VVol, regardless of any sharing or thin-provisioning considerations. This sizing is true of all types of VVols, including clones and snapshots

The array implementation uses a 10x multiplier of the storage container physical size when computing the storage container's logical size. Therefore, a storage container with physical size of 10GB will have a logical size of 100GB.

Create a Storage Container

NOTE: You must be logged in as group administrator to create a storage container.

To create a new storage container:

1. Click the **VMware** tab or icon.
2. Select **Storage Containers**.
3. In the Activities panel, click **Create storage container**.

   The Create Storage Container wizard opens.
4. On the **General settings** page:
   a. (Required) Type a name for the container.
   b. (Optional) Type a description for the container.
   c. If more than one storage pool is listed in the table, select the storage pool to be used for the container.
5. Click **Next**.
6. On the **Physical space** page:
   a. (Required) Type a value in the **Container physical size** field.
   b. Select the unit of measurement (MB, GB, TB) for the physical size.
7. Click **Next**.
8. On the **Summary** page, review the settings, then click **Finish** to create the storage container.

Modify a Storage Container

NOTE: Storage containers cannot be explicitly moved between pools. If two pools are merged via either the “pool merge” or “pool delete” operation, any storage containers that draw from the deleted or merged pool are automatically reassigned to the default or the merged-to pool respectively.

To modify attributes of an existing storage container:

1. Click the **VMware** tab or icon.
2. Select **Storage Containers**.
3. In the **Storage Containers** panel, select the container that you want to modify.

4. In the Activities panel, click **Modify storage container**.

   The Modify Storage Container dialog box opens.
5. On the **General settings** tab:
   a. Modify the name for the container and/or the description.
   b. Select a different storage pool for the container (if more than one pool is available).
6. On the **Physical space** tab:
   a. Type a value in the **Container physical size** field.
   b. Change the unit of measurement for the size (MB, GB, TB) if necessary.
7. Click **OK** to modify the storage container.

Delete a Storage Container

NOTE: You cannot delete a storage container if it contains VVols.

To delete a storage container:

1. Click the **VMware** tab or icon.
2. Select **Storage Containers**.
3. In the **Storage Containers** panel, select the container that you want to delete.
4. In the Activities panel, click **Delete storage container**.
5. In the Delete Storage Container confirmation dialog box, click **Yes** to delete the storage container.

Virtual Machines Tab

To view the virtual machines in your PS Series Group, click **VMware**, then under the **VMware** panel select **Virtual machines**.

Table 38. Information Identified in Virtual Machines Tab identifies the information shown for virtual machines.

Table 38. Information Identified in Virtual Machines Tab

Column	Description
Virtual machines	
VM	Name of the virtual machine
Guest OS	Guest operating system running on the virtual machine
Allocated space	Amount of space allocated for the virtual machine
Description	Text string that describes the virtual machine
VVols	
VVol	Identification of the VVol
Type	VVol types, such as: config, data, swap, memory
Storage container	Name of the storage container
Reported size	Size reported for the virtual disk
Status	Status (bound or unbound)
Last bound time	Date and time when the last bind operation occurred
Snapshots	Number of snapshots of the virtual disk
Linked clones	Number of linked clones of the virtual disk

NOTE: The Dell Virtual Storage Manager (VSM) runs an event process that enables Group Manager to show the correlation between the virtual machine and the VVol. In some cases, the vCenter virtual machine-related event might be incorrectly propagated to Group Manager. To manually trigger a verification between the VSM and Group Manager, use the VSM system schedule command Verify Snapshots and Replicas Run Now. For details, see the *Dell Virtual Storage Manager Installation and User’s Guide*.

11

NAS Operations

Network-attached storage (NAS) provides high-performance, high-availability, scalable resources with on-demand provisioning in a unified storage environment.

You can perform basic and advanced operations on NAS storage, as shown in Table 39. Basic and Advanced NAS Operations.

Table 39. Basic and Advanced NAS Operations

	NAS Cluster Operations	NAS Controller Operations	NAS Container Operations
Basic	Configure a NAS cluster Modify NAS cluster settings Add, modify, or delete a local group for a NAS cluster	Add or replace NAS controllers Update NAS controller firmware	Create, modify, or delete NAS containers Create, modify, or delete an SMB share Create, show, or delete an SMB home share Create, modify, or delete an NFS export
Advanced	Diagnose problems in a NAS cluster Delete a NAS cluster	Shut down a NAS controller pair	Rebalance SMB client connections across NAS containers Add, modify, or delete a NAS antivirus server Protect NAS container data with NDMP Modify the file security style for NAS containers Create, modify, or delete quotas Enable thin provisioning and data reduction Reinstall the FluidFS operating system

12

NAS Cluster Operations

Table 40. Basic and Advanced NAS Cluster Operations provides a list of basic and advanced NAS cluster operations.

Table 40. Basic and Advanced NAS Cluster Operations

Basic	Configure a NAS cluster Modify NAS cluster settings Add, modify, or delete a local group for a NAS cluster
Advanced	Diagnose problems in a NAS cluster Delete a NAS cluster

NAS Cluster Configuration

A NAS cluster is a collection of NAS appliances configured in a PS Series group. Configure NAS clusters as follows:

- 1 appliance with 2 NAS controllers, in one of the following configurations:
  - 2 Model FS7500
  - 1 Model FS7600
  - 1 Model FS7610
- 2 appliances with 4 NAS controllers (maximum configuration), in one of the following configurations:
  - 4 Model FS7500
  - 2 Model FS7600
  - 2 Model FS7610
  - 2 Model FS7500 and 1 Model FS7600 (mixed cluster)

For each controller in the NAS cluster, you need IPv4 addresses for each of the following networks:

- Client network
- SAN network

NOTE: The actual number of IPv4 addresses varies depending on your system configuration.

During the initial configuration, you need only one NAS cluster IPv4 address to configure the NAS cluster.

- FS7500 and FS7600 NAS clusters support four IPv4 addresses per controller.
- FS7610 NAS clusters support two IPv4 addresses per controller.

About Mixed-Model NAS Clusters

Mixed NAS clusters contain different NAS appliance models. You can create mixed NAS clusters during NAS cluster configuration, or you can expand an existing NAS cluster by adding a NAS controller pair, as described in Table 41. NAS Cluster Expansion.

Table 41. NAS Cluster Expansion

Existing NAS Cluster	Expansion NAS Controllers
Two model FS7500 NAS controllers	Two FS7500
	One FS7600 (creates a mixed cluster)
One model FS7600 NAS appliance	One FS7600
	Two FS7500 (creates a mixed cluster)
One model FS7610 NAS appliance	One FS7610

Configure a NAS Cluster

You must configure two NAS controllers to a NAS cluster as a NAS controller pair; you cannot add just one NAS controller. However, you can detach and attach a single NAS controller to a NAS cluster if you need to replace a NAS controller.

NOTE: The first time you configure a NAS cluster, you are asked to accept an End-User License Agreement (EULA). You cannot configure the NAS cluster unless you accept the EULA.

Prerequisites

Before you configure a NAS cluster:

- Satisfy network requirements:
  - Enable jumbo frames on the switch ports used for the internal and SAN networks.
  - Enable flow control on the same ports.
- Verify that a storage pool in the PS Series group has sufficient space to store the NAS reserve and client data. The NAS reserve contains NAS client data available for use by NAS containers, and internal capacity reserved for system use.

  After you configure the NAS cluster, you can increase the size of the NAS reserve as needed.

  NOTE: You cannot decrease the size of the NAS reserve.
- Verify that you have the following information:
  - Name to assign to this NAS cluster
  - Client network information
  - SAN network information
  - Password for the grpadmin account

Enter your NAS cluster settings in Table 42. NAS Cluster Configuration Settings.

Table 42. NAS Cluster Configuration Settings

Name of NAS cluster:		
Client network		
	Client network name	
	IP address	
	Netmask	
	Default gateway	
	NAS controller IP addresses	

SAN network		
	NAS cluster management IP address	
	NAS controller IP addresses	
Password for grpadmin account:		

NOTE:

- If you are using DNS, you must manually add the NAS cluster IP address and NAS cluster name to the DNS server.
- If you are in a routed client network and using multiple NAS cluster IP addresses, add all NAS cluster IP addresses to the DNS server and associate them with the same NAS cluster name.

NAS Cluster Configuration

1. Click **NAS**, and when the message is displayed that prompts you to configure a NAS cluster, click **Yes**.

   After the NAS devices have been discovered, the available NAS controllers appear in the Discover NAS Devices dialog box.
2. Select the controllers that make up the cluster. FS7500 controllers must be selected in pairs. FS7600 controller pairs are collocated in a single enclosure, called a NAS appliance, which you select with a single checkbox.
3. Click **Configure NAS cluster**. The End-User License Agreement appears.
4. Select the **I accept the terms of the End-User License Agreement** checkbox and click **OK**.
5. On the Getting Started window, click **Next**.
6. In the Client Network window, specify:

   - Name

     NOTE: The NAS cluster name can contain up to 15 ASCII letters, numbers, and hyphens.
   - IP address
   - Netmask
   - Default gateway
7. In **NAS Controller IP addresses for the client network**, type one IP address for each NAS controller. Specify a minimum of one and a maximum of eight NAS cluster IP addresses. (You can also click **Auto fill** to automatically enter the NAS Controller IP addresses based on the first NAS cluster IP address.) Click **OK**.
8. Under **SAN Network Settings**, type the NAS cluster management IP address. The **Group IP Address** and **Netmask** fields are preset based on the group configuration.
9. Under **NAS Controller IP addresses for the SAN network**, click the **Auto fill** button to populate the table with NAS controller SAN IP addresses and click **Next**.

   NOTE: The SAN network and internal network must be in the same VLAN.
10. Review your parameters on the Summary screen and click **Finish** when you are ready to proceed. The clusterization process begins running.

If all configuration parameters are valid, all the steps in the wizard display a green checkmark. If not, a window lists the invalid parameters with red Xs. A new Change Configuration button appears to the left of the Retry button. This button allows you to make corrections to the invalid parameters.

Before clicking the **Change Configuration** button, verify that:

- NAS controllers are correctly connected to the appropriate switch stack (that is internal, SAN, and IPMI ports are connected to one switch stack, and client ports are connected to the client switch stack).
- No duplicate IP addresses are in the network configuration. The IP addresses used in the NAS configuration and the group configuration must be unique in the network. If an IP address used in the NAS cluster is also used elsewhere, change the IP address where it is used outside of the NAS cluster.

Change the invalid configuration parameters. Clicking **OK** will automatically retry the specified configuration. The controller statuses are reset to in progress and the red Xs are cleared from the display. If the retry operation was successful, the status shows success.

Or, you can click **Retry**. If the network configuration fails again, contact your service representative.

The group remains operational during the NAS cluster configuration. When the configuration succeeds, the progress window displays a success message.

NOTE: During the NAS cluster configuration, you might see the following message:

```
Status of 0 is failed. Status of 1 is failed.
```

You can disregard this message.

Resolve a Failed NAS Configuration

If a configure NAS cluster operation fails, you can consider one of the solutions described in the following sections. If you cannot correct the problem, contact your customer support representative.

If an operation fails during a NAS cluster configuration, Group Manager gives you the opportunity to correct the problem. You can then click **Retry** to retry the failed operation and resume the NAS cluster configuration.

In addition, some configuration problems can be ignored. When they occur, Group Manager displays a message indicating the problem. You can then click **Continue** to resume the NAS cluster configuration without retrying the failed operation. See your NAS appliance hardware documentation for information about determining hardware status and correct the problem using the following procedures.

Member Configuration Failure

Description: You successfully discovered the NAS controllers, specified the NAS cluster configuration, and started the configuration operation, but the NAS controller configuration fails.

Verify that:

- NAS controllers' network ports are correctly connected to the proper client and SAN networks.
- Make sure that you are the only PS Series group administrator trying to configure the NAS controllers into a NAS cluster.

Click **Retry**. If the NAS controller configuration fails again, contact your customer support representative.

Network Configuration Failure

Description: Failure occurs during SAN or client network configuration.

Verify that:

- Network requirements are met:
  - Enable jumbo frames on the switch ports used for the internal and SAN networks.
  - Enable flow control on the same ports.
- NAS controllers are correctly connected to the appropriate switch stack (SAN and IPMI ports are connected to the SAN switch stack, and client ports are connected to the client switch stack).
- No duplicate IP addresses are in the network configuration. The IP addresses used in the NAS configuration and the group configuration must be unique in the network. If an IP address used in the NAS cluster is also used elsewhere, change the IP address where it is used outside of the NAS cluster.

Click **Change Configuration**.

Depending on the type of configuration that failed, you will see one of the following pages:

- If the SAN network failed, the **Modify SAN network** page opens.
- If the client network failed, the **Modify client network** page opens.

You can change any of the settings that were not already stored in the system on these pages. Click **OK** to accept your changes.

Click **Retry**. If the network configuration fails again, contact your customer support representative.

Validation Failure

Description: You see a `system validation error` or similar error message.

NOTE: If you are using the minimum network configuration (not supported for production environments), ignore this error and click Continue to resume the NAS cluster configuration.

If you are using the recommended network configuration, click the error message to display details. For example:

- Disconnected cable

  Verify the cable connections for your hardware configuration.
- HW mismatch error

  Contact your customer service representative.
- BPS error

  Check the BPS connectivity to the power grid and to the NAS controllers. Then, click **Retry**. If the validation fails again, contact your customer support representative.
- IPMI error

  NOTE: For the FluidFS 7500 only, ensure that the IPMI port (lower-left corner of the NAS controller back panel) is connected to the SAN switch stack. Then, click Retry. If the validation fails again, contact your customer support representative.
- Network error

  Follow the recommendations above for handling a network configuration failure and ensure that the switches are configured for jumbo frames.

Click **Retry**. If the validation fails again, contact your customer support representative.

Format Failure

If you experience a failure during the format operation:

1. Click **Volumes** in the lower-left corner of the Group Manager window.
2. Expand **Volumes** in the tree, select a NAS volume (the NAS volume name is prepended with the NAS cluster name), and then click the **Access** tab.
3. Repeat this task for each NAS volume. Each NAS controller should have one or four iSCSI connections to each NAS volume, depending on the number of SAN access IP addresses that you specified for each NAS controller.
4. If you do not see the required number of connections, check all your network connections.
5. Click **Retry**. If the format operation fails again, contact your customer support representative.

NAS Cluster Post-Setup Tasks

After you configure a NAS cluster, perform these tasks:

- Review NAS cluster information.

  This information includes status, space utilization, member information, and network configuration.
- Add NAS cluster IP addresses.

  If the NAS cluster client network is a routed network, you can modify the client network configuration and add more NAS cluster IP addresses for proper load balancing across multiple client subnets.
- Configure local users and groups to control access to SMB shares, SMB home shares, and NFS exports.

  The NAS cluster authenticates local users and groups.

- Configure the NAS cluster to use DNS (Domain Name Service), which is a networking service that translates Internet domain names into IP addresses.

  If you want to use DNS, manually add the NAS cluster IP address and NAS cluster name to the DNS server. If you are in a routed client network and using multiple NAS cluster IP addresses, add all the NAS cluster IP addresses to the DNS server and associate them with the same NAS cluster name.

  In addition, you must configure DNS to use Active Directory in a NAS cluster.
- Modify the default values for NAS container space and permission settings.

  When a file or directory is created, the default NAS container security style, which controls the permissions and ownership, is applied.

  For a new NAS container, you can modify clusterwide default values:

  - NAS container data reduction settings, including adding compression and defining policies
  - NAS container in-use space warning limit, snapshot reserve percentage, and snapshot in-use space warning limit
  - File security style (Mixed, NTFS, or UNIX)
  - UNIX permissions for files and directories created in the NAS container by Windows clients using SMB (only for the UNIX file security style)

NOTE: Select the Enable data reduction checkbox to activate the Modify policy button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.

- Modify the default values for NFS export permission and trusted user settings.

  When you create an NFS export, the service applies default values for the permission and trusted user settings. You can modify the following NAS clusterwide default values for a new NFS export:

  - Read-write or read-only permission
  - Trusted users (All except root, all, or nobody)
- Configure external authentication in the NAS cluster to control access to SMB shares, SMB home shares, and NFS exports.

  You can configure Active Directory for authenticating Windows users and groups, and you can configure NIS or LDAP for authenticating UNIX users and groups. You can also create mappings between Windows and UNIX users.
- (NAS) Create a NAS container. You can create multiple NAS containers in a NAS cluster.

After you create a NAS container, perform these tasks:

- Display NAS container information.

  You can display information about the NAS containers in the NAS cluster, including the status, space utilization, SMB shares, NFS exports, snapshots, schedules, and quotas.
- Modify the file access security style.

  You can modify the file access security style (Mixed, NTFS, or UNIX) for a NAS container.
- Modify the UNIX permissions for Windows files and UNIX permissions for Windows directories.

  If the NAS container is using the UNIX file security style, you can modify the UNIX file and directory permissions (Read, Write, and Execute) for Owner, Group, and Others.
- Create quotas.

  You can create group and user quotas to control client space usage in a NAS container.
- Create an SMB share.

  You can create multiple SMB shares on a NAS container.
- Enable an SMB home share.

  An SMB share is not required to enable SMB home shares, all you need is a NAS container.
- Set the SMB administrator password.

  You must set the password to access SMB shares.
- Create an NFS export.

  You can create multiple NFS exports on a NAS container.

- Create a snapshot.

  To protect NAS container data, you can create snapshots.
- Create a snapshot schedule.

  Use a schedule to regularly create NAS container snapshots.
- Configure NAS container replication.

Modify a NAS Cluster Name

To modify the name of a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Rename NAS cluster**.
3. In the **Rename NAS cluster** dialog box, specify the new NAS cluster name. The NAS cluster name can contain up to 15 ASCII letters, numbers, and hyphens.
4. Click **OK**.

   NOTE: The NAS cluster name can contain up to 15 ASCII letters, numbers, and hyphens.

Modify NAS Clusterwide Default NAS Container Settings

To modify or display the NAS clusterwide default NAS container space settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default NAS Container Settings** panel, modify the settings as needed.

Select an NFS Protocol Version

You can choose among the three supported NFS protocol versions (NFSv3, NFSv4.0, or NFSv4.1) of the NAS System for the NFS mode.

To select an NFS protocol version:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab.
3. In the NFS panel, select an NFS version from the **Maximum Supported NFS Protocol Version** drop down menu.

Modify the Size of the NAS Reserve

NOTE: You cannot decrease the size of the NAS reserve.

To increase the size of the NAS reserve:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Modify NAS reserve**.
3. In the NAS reserve section, specify the new Total Space or Usable Space size and click **OK**.

A progress window opens in the temporary Operations tab, showing the progress of the resize NAS reserve operation. This operation might take up to 30 minutes to complete.

NOTE: Wait for the NAS reserve resize operation to complete before starting another long-running operation, such as adding a NAS controller pair, detaching a NAS controller, or attaching a NAS controller operation.

If the NAS reserve resize operation fails, contact your Dell support representative.

Add a Local Group for a NAS Cluster

To add a local group for a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Local Users and Groups** tab.
3. In the Local Groups panel, click **Add**.
4. In the **Add local user** dialog box, specify the group name.

   NOTE: The group name accepts up to 20 ASCII characters, including letters, numbers, underscores, hyphens, and periods. The first character must be a letter or a number.
5. Click **OK**.

   NOTE: You cannot modify a local group. Instead, you must delete the local group and then add a new group.

Delete a Local Group from a NAS Cluster

To delete a local group:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Local Users and Groups** tab.
3. In the Local Groups panel, select the group and click **Delete**.

   NOTE: You cannot delete the Administrators group.
4. When prompted, confirm that you want to delete the local group.

Add a Local User on a NAS Cluster

To add or modify a local user for a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Local Users and Groups** tab.
3. To add a local user, in the Local Users panel click **Add** and specify the user information.

   NOTE:

   - **Local user names can accept up to 20 ASCII characters, including letters, numbers, underscores, hyphens, and periods, but cannot end with a period. The first character must be a letter or a number.**
   - **Local user passwords must contain between 7 and 20 ASCII characters.**
   - **Local user descriptions (optional) can accept up to 63 Unicode characters. The description is truncated if you enter more than 63 characters. (Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.**)
4. Click **OK**.

Modify a Local User on a NAS Cluster

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Local Users and Groups** tab.
3. In the Local Users panel, click **Modify** and specify the user information.

   NOTE:

   - **Local user names can have up to 20 ASCII characters, including letters, numbers, underscores, hyphens, and periods, but cannot end with a period. The first character must be a letter or a number.**
   - **Local user passwords must contain between 7 and 20 ASCII characters.**
4. Click **OK**.

Delete a Local User from a NAS Cluster

To delete a local user from a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Local Users and Groups** tab.

3. In the Local Users panel, select the user and click **Delete**.

   NOTE: You cannot delete the built-in local administrator account (Administrator).

4. Confirm that you want to delete the local user.

Map Users for a NAS Cluster

NOTE: To map users, you must have Active Directory and either LDAP or NIS configured in the NAS cluster.

To define a mapping between a Windows user and a UNIX user:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the User Mapping panel, click **Map users**. The Manual User Mapping dialog box opens.
4. To add a user mapping, click **Add**.
5. Type the name or partial name of the UNIX or Windows user that you want to search for and click **Search**.

   NOTE: You can search for a Windows user in a specified subdomain by selecting the subdomain from the Search Subdomain drop-down menu.

6. In the Manual Map Users dialog box, select the UNIX user and the Windows user. Then, open the **Direction** menu and select the user-mapping direction:
   - Windows user to the UNIX user
   - UNIX user to the Windows user
7. Click **OK** to add the new user mapping. The Manual User Mapping dialog box opens with the new mapping.
8. Click **OK** in the Manual User Mapping dialog box.

Set the User Mapping Policy for a NAS Cluster

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the User Mapping panel, either:
   - Select whether to automatically map all Windows users in Active Directory to the identical UNIX users in NIS or LDAP. The default setting is no automatic mapping.
   - Click **Map users** to manually map users.

NOTE: You can use the Group Manager CLI to specify whether to map unmapped users to the guest account (the default setting). See the *Dell EqualLogic Group Manager CLI Reference Guide*.

Delete a User Mapping for a NAS Cluster

To delete a user mapping:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the User Mapping panel, click **Map users**. The Manual User Mapping dialog box opens, displaying the current user mappings.
4. Select the user mapping and click **Delete**.
5. Confirm that you want to delete the user mapping.

Configure an Active Directory for a NAS Cluster

Add a NAS cluster to Active Directory before you can use Active Directory to externally authenticate NAS cluster users.

Prerequisites:

- The Active Directory server and the PS Series group must use a common source of time.
- Configure the NAS cluster to use DNS. The DNS servers that you specify must be the same DNS servers that your Active Directory domain controllers use.

You can specify an NTP server for the group and a DNS server for the NAS cluster when configuring Active Directory, or you can perform these tasks separately.

NOTE: Configuring Active Directory interrupts client access to SMB shares.

To configure Active Directory:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the Active Directory panel, click **Configure Active Directory**. The Getting Started window of the Configure Active Directory wizard opens.
4. Click **Next**. The Configure NTP Servers window opens.
5. (Optional) Click **Add**, enter the IP address of the NTP server and the port to use (the default port is 123), and click **OK**. You can add up to three IP addresses for NTP servers.
6. Click **Next**. The Configure DNS window opens.
7. In the DNS servers list box, in order of preference, click **Add** to specify a DNS server. You can define up to three servers. The servers are used in the order in which they are listed. Use the **Up** and **Down** links to organize the servers in the preferred order.
8. In the DNS suffixes list box, in order of preference, click **Add** to specify a DNS suffix. You can specify up to three suffixes. The suffixes are used in the order in which they are listed. Use the **Up** and **Down** links to organize the servers in the preferred order.
9. Click **Next**. The Configure Active Directory window opens. Enter the Active Directory configuration information.
   a. You can optionally configure preferred domain controllers in this window. On the lower portion of the Active Directory window, select the **Configure preferred domain controllers** checkbox.
   b. Select a domain controller and click **Add** or **Delete**.
10. If the Active Directory configuration is correct, click **Finish** in the final dialog box to complete the configuration. Click **Back** to make changes.

After Active Directory is configured successfully, the NAS cluster joins the Active Directory.

NOTE:

- **When you configure an Active Directory with preferred domain controller join, if an error message is generated, the domain controller is not added and must be added manually.**
- **You cannot delete the Active Directory configuration. To stop using Active Directory, click Leave.**

Configure Preferred Domain Controllers

Setting a preferred DC enables you to limit traffic to one DC or another, use a DC based on physical location, or specify a different DC if your primary or preferred DC is not available. If the DC list is modified, the preferred DC list is automatically updated so that you do not need to leave AD and rejoin.

To configure preferred domain controllers:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the Active Directory panel, click **Configure preferred domain controllers**.

   NOTE: To see this option, Active Directory must already be configured.
4. In the Configure Preferred Domain Controllers dialog box, click **Add**.
5. Type the host name or IP address of the preferred domain controller, and click **OK**.
6. Click **OK**.

Leave Active Directory

To leave an Active Directory domain:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab. The status of Active Directory is `enabled`.

3. In the Active Directory panel, click **Leave**.

   Domain users will be prevented from access if you leave the Active Directory domain. The status changes to `not configured`.

   NOTE: You cannot delete the Active Directory configuration.

Configure or Modify NIS or LDAP for a NAS Cluster

To authenticate UNIX clients, you can use NIS or LDAP for external authentication.

NOTE: Configuring NIS or LDAP interrupts client access to SMB shares.

To configure NIS or LDAP in a NAS cluster or to modify the existing configuration:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the UNIX Authentication panel, click **Configure UNIX authentication**.
4. In the UNIX Authentication Server Configuration dialog box, select **NIS** or **LDAP** and enter or modify the NIS or LDAP information.
5. Click **OK**.

Delete NIS or LDAP Configuration for a NAS Cluster

You can delete the current NIS or LDAP configuration in the NAS cluster. If you want to resume use of NIS or LDAP, you must configure NIS or LDAP again.

To delete the NIS or LDAP configuration:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. In the UNIX Authentication panel, click **Remove UNIX authentication**.
4. Confirm that you want to remove the UNIX authentication.

Modify the Client Network Configuration

To modify the client network configuration:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Network** tab, then double-click the client network to be modified. The Modify Client Network dialog box opens.
3. Modify the following fields as needed:

   - Name

     The client network name is **Client** by default and cannot be changed.
   - VLAN tagging

     Configure ID numbers ranging from 1–4094. The default VLAN is 0.
   - Netmask

     Can be modified.
   - Virtual IP addresses

     Click **Auto fill** to automatically enter NAS cluster IP addresses based on the first IP address in the above list. If you delete an IP address, any client connected to the IP address is disconnected. The client usually reconnects automatically.
   - NAS controller IP addresses for the client network

     Click **Auto fill** to automatically enter NAS controller IP addresses based on the first NAS cluster IP address. If you change a NAS controller IP address, clients might be disconnected. The client usually reconnects automatically.

4. Click **OK** to apply the changes.
5. In the Activities panel, click **Modify client properties** to open the dialog box.
6. Modify the following properties as needed:

   - Default gateway
   - MTU size

     If you change the MTU size, clients are disconnected. The client usually reconnects automatically.

     NOTE: Modify the MTU byte size only if directed by Dell support. For normal NAS cluster operation, a value of 1500 is required.
   - Bonding mode (ALB or LACP)
7. Click **OK** to apply the changes.

After you modify the client network configuration, you must change the primary client subnet.
To change the subnet:

1. Delete the default gateway by setting it to 0.0.0.0.
2. Change the subnet.
3. Configure a new default gateway for the new subnet.

Configure DNS for a NAS Cluster

Domain Name Service (DNS) is a networking service that translates Internet domain names into IP addresses.

If you want to use DNS, manually add the NAS cluster IP address and NAS cluster name to the DNS server. If you are in a routed client network and using multiple NAS cluster IP addresses, add all the NAS cluster IP addresses to the DNS server and associate them with the same NAS cluster name.

In addition, you must configure the NAS cluster to use DNS in order to use Active Directory in a NAS cluster. The DNS servers that you specify must be the same DNS servers that your Active Directory domain controllers use.

To configure DNS for a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Configure DNS**.
3. In the **Configure DNS** dialog box:

   a. Specify one to three IP addresses for DNS servers. The order in which you specify the servers is the order in which they are used.

   b. Specify one to three DNS suffixes (for example, company.com).

      NOTE: A DNS suffix accepts up to 126 ASCII characters: letters, numbers, period, hyphen, and underscore.
4. Click **OK**.

About the Internal Network Required for NAS Configuration

The internal network is used for communication between NAS controllers. A NAS cluster requires a block of IP addresses for the internal network. Starting with FS Series firmware v4, the internal network is automatically configured for you. Internal network-related information is no longer displayed to users in Group Manager or the CLI.

SAN Network Configuration Settings

The network that is used for access between the PS Series group (SAN) and the NAS controllers is typically separate from the client network.

The SAN network configuration for a NAS cluster includes the following IP addresses:

- Management IP address, which allows access between the PS Series group and the NAS cluster. The management IP address must be on the same subnet as the group IP address.
- IP addresses for each NAS controller, which allows access between the PS Series group and the NAS controllers. The NAS controller IP addresses must be on the same subnet as the group IP address. You must specify the maximum IP addresses for your NAS appliance hardware IP addresses for each NAS controller.

  The number of SAN access IP addresses must be the same on each NAS controller in the NAS cluster.

By default, the SAN network uses 9000 (jumbo frames) for a maximum transmission unit (MTU) size. You cannot change the MTU size for the SAN network.

If you modify the PS Series group network configuration (for example, if you modify the group IP address), you might need to make a reciprocal adjustment to the NAS cluster SAN network configuration, including the SAN netmask. The SAN network configuration must be in the same subnet as the group IP address.

Before you can modify the SAN network configuration, you must put the NAS cluster into maintenance mode.

NOTE: After you perform a SAN network change with replication configured, you will need to manually create a new static route for the new SAN network.

Modify the SAN Network Configuration

If you modify the PS Series group network configuration (for example, if you modify the group IP address), you might need to make a reciprocal adjustment to the NAS cluster SAN network configuration.

NOTE: Before you can modify the SAN network configuration for a NAS cluster, you must put the NAS cluster into maintenance mode.

To modify the SAN network configuration:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Modify SAN network**.
3. In the **Modify SAN network** dialog box, you can modify:
   - Management IP address
   - SAN netmask (based on the group IP address subnet)
   - IP addresses for each NAS controller (either 1 or 4)

     The number of SAN access IP addresses for each NAS controller must be the same on each NAS member (NAS controller pair). Click the **Auto Fill** button to automatically enter IP addresses based on the next available addresses after the NAS cluster management IP address.

   NOTE:

   - **The SAN network configuration must be in the same subnet as the group IP address.**
   - **After you perform a SAN network change with replication configured, you will need to manually create a new static route for the new SAN network.**
4. Click **OK**.

About NAS Cluster Maintenance Mode

Maintenance mode is a tool for the system administrator to take the file system offline, thereby isolating the NAS cluster from outside connections and file-system activity. The main purpose of maintenance mode is to provide the administrator with a safe and organized method to take down a file system without harming or losing current data, before shutting down the controllers or storage array when it is necessary to do so.

While the NAS cluster is in maintenance mode, all client connections to the NAS controllers are stopped and no new connections can be made. Put the NAS cluster into maintenance mode when you:

- Move the NAS hardware to a different location.
- Change the PS Series group IP address because of:
  - Network changes in your environment
  - A new configuration, such as adding a management network to the group
- Perform maintenance or infrastructure work.
- Want to isolate the system in order to address specific issues, such as:
  - Analyze attack attempts
  - Resolve traffic issues

When you have completed any of these tasks in maintenance mode, set the NAS cluster back to normal mode.

Set a NAS Cluster to Maintenance Mode

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. In the Activities panel, click **Enter Maintenance Mode**.
3. Confirm the operation.

While the cluster is changing to maintenance mode, the Activities link for **Enter maintenance mode** becomes dimmed and the cluster status reports `Transitioning to maintenance`. When the transition is complete, the Activities link changes to **Enter normal mode**.

Resume Normal Cluster Operation

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Enter normal mode**.

Shut Down and Restart a NAS Cluster Manually

In rare circumstances, you might need to shut down a NAS cluster. For example, you might need to shut down the cluster if you are moving the NAS hardware to a different location.

To shut down a NAS cluster:

1. Inform clients that the NAS cluster will be unavailable for some time.
2. Put the NAS cluster into maintenance mode.
3. Use the Group Manager GUI to identify the service tag identification numbers for the controller pair.
4. See your hardware documentation to:
   a. Determine how to find the service tag information for a specific NAS appliance. This number might be in the controller firmware (software accessible) or on a physical label or tag on the NAS appliance chassis.
   b. Find the instructions for shutting down a NAS controller.
5. Power off each NAS controller in turn. See one of the following manuals:
   - *Dell Equallogic FS7500 Series Appliances Hardware Owner's Manual*
   - *Dell Equallogic FS7600 Series Appliances Hardware Owner's Manual*
   - *Dell Equallogic FS7610 Series Appliances Hardware Owner's Manual*

To restart NAS cluster operation after you reinstall the hardware:

1. See your hardware documentation to find the instructions for applying power to a NAS controller.
2. Apply power to each NAS controller used in the NAS cluster.

3. After the hardware restarts, start the NAS cluster.

About Deleting a NAS Cluster

If you no longer need to provide NAS operations, you can delete the NAS cluster from the PS Series group.

When you delete a NAS cluster:

- All service data and all client data that is stored in the NAS reserve is destroyed.
- The NAS reserve space is added back to the free pool space.
- The NAS controllers are reset to the factory defaults.
- The NAS controllers reboot.

When you delete the NAS cluster, a message appears stating that the cluster was deleted successfully. However, you cannot discover the NAS controllers and use them in a NAS cluster until the NAS controllers reboot.

You can delete a NAS cluster if some NAS controllers are down. If any controllers are down, you should confirm that you want to continue.

You must reinstall the firmware to reset NAS controllers to their factory defaults.

Delete a NAS Cluster

Deleting a NAS cluster destroys all NAS cluster configuration data and all client data that is stored in the NAS reserve, and the NAS reserve space becomes free pool space. The NAS controllers are reset to the factory defaults, and will reboot.

To delete a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click **Delete NAS cluster**.
3. Select the **Destroy *size* data** checkbox, where *size* is the size of the NAS cluster that you want to delete. Select this checkbox to delete the cluster.
4. Click **Yes**.

   A message is displayed stating that the NAS cluster was deleted successfully. However, you must wait until the NAS controllers reboot before you can discover them.

13

NAS Controller Operations

Table 43. Basic and Advanced NAS Controller Operations provides a list of basic and advanced NAS controller operations.

Table 43. Basic and Advanced NAS Controller Operations

Basic	Add or replace NAS controllers Update NAS controller firmware
Advanced	Shut down a NAS controller pair

Add Additional NAS Controllers

After you add NAS controllers using the wizard, you can improve performance and availability of your network by adding up to two NAS appliances to the NAS cluster.

NOTE: While you are adding a NAS controller pair, you cannot perform any other NAS cluster operation until the add controller pair operation completes.

Before adding a NAS controller pair, see your hardware documentation to obtain the following information and verify the state of the NAS controller:

- The NAS controllers are connected to the networks and power.
- Each NAS controller is powered on and is in standby mode.
- Obtain the service tag identification number for each controller that you are adding.

Obtain and verify the following network resources:

- 500GB of free NAS reserve space in the PS Series group.
- 1 client network IP address for each NAS controller in the controller pair.
- 4 or 8 SAN access IP addresses for each NAS controller in the NAS controller pair. The number of SAN access IP addresses for each new NAS controller must be the same as for the existing NAS controllers.

To add a NAS controller pair:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click **Add NAS controller pair**. (This link appears only if the group already contains a pair of NAS controllers.)
3. Click **Yes** to confirm that you want to add a controller pair.
4. After the NAS devices have been discovered, the available NAS controllers appear in the **Add NAS controller pair** dialog box. In the list of discovered devices, select the service tag numbers of the NAS controllers that you are adding.

   FS7600 or FS7610 controller pairs share a common service tag identification:

   - *number*-1
   - *number*-2
5. Click **Add NAS controller pair**. The Getting Started tab of the Add NAS Controller Pair wizard opens.

   The **Add NAS controller pair** button is not enabled until you have selected a valid pair of NAS controllers.
6. Click **Next**. The Client Network tab opens.

7. For each client network, verify the following settings for the NAS cluster, and click **Auto fill**:

   - VLAN tagging
   - IP address
   - Netmask
   - Default gateway

8. Click **Auto fill** to populate the table with NAS controller IP addresses, or type the addresses.

   NOTE:

   - **The Auto Fill option bases new addresses on the first NAS cluster IP address. This approach results in duplicate addresses if any of the new addresses are already used on the network. You might want to specify IP addresses individually to avoid duplicate addresses.**
   - **Repeat this procedure for each configured client network. You will not be able to proceed until the information is filled in for each client network.**

9. Click **Next**. The SAN Network tab opens.
10. Verify the following SAN network information:

    - Group IP address
    - NAS cluster management IP address
    - Netmask

11. Click **Auto fill** to populate the table with IP addresses for the SAN network, or type the addresses.

    NOTE: The Auto Fill option bases new addresses on the first NAS cluster IP address. This approach results in duplicate addresses if any of the new addresses are already used on the network. You might want to specify IP addresses individually to avoid duplicate addresses.

12. Click **Next**. The Summary tab opens.
13. In the Summary tab, either:

    - Click the **Copy** link and save the summary information in a text file.
    - Click **Back** to return to a previous page to verify the NAS controller information.
    - Click **Finish** to add the NAS controller pair.

About Replacing a NAS Controller with Detach and Attach

In rare cases, you must remove a NAS controller from a NAS cluster and replace it with a different NAS controller. For example, you might need to replace a malfunctioning NAS controller.

You can remove a NAS controller from a NAS cluster without affecting data availability by detaching the controller. After you detach the controller, it resets to its factory defaults and powers off, if possible.

While a NAS controller is detached from the NAS cluster, SMB shares and NFS exports remain available. Most NAS cluster configuration changes are not allowed while a NAS controller is detached.

To resume full NAS cluster functionality, you must attach a functioning NAS controller to the service. After the new controller is attached, it inherits the NAS cluster configuration settings of the original NAS controller.

NOTE: Detach a NAS controller only under the direction of your customer support representative and only when you are replacing that NAS controller.

Some NAS controller maintenance operations do not affect NAS controller operation (for example, replacing a hot-swappable power supply). Other NAS controller maintenance operations require you to cleanly shut down the NAS controller, perform the maintenance, and then restart the NAS controller. See your hardware documentation.

Detach a NAS Controller

To replace a malfunctioning NAS controller without affecting data availability, you can detach the NAS controller and then attach a functioning NAS controller. A detach operation should be used only to replace a NAS controller.

NOTE: Detach a NAS controller only when directed by your customer support representative.

In some cases, your support representative might instruct you to cleanly shut down a NAS controller before detaching it. After cleanly shutting down the NAS controller, you can turn on power, wait for the NAS cluster to recognize the NAS controller, and then detach the NAS controller.

You cannot detach a NAS controller if its peer NAS controller is already detached. However, in a 4-controller NAS cluster, you can detach one controller in each NAS controller pair.

To detach a NAS controller:

1. Click **Group Configuration**, then expand **Members**.
2. Select a NAS member and, if necessary, expand it to show the NAS controller name.
3. After verifying the status of its peer NAS controller, select an attached NAS controller.
4. Click **Detach NAS controller**.
5. Confirm that you want to detach the NAS controller.

A progress window opens, showing the progress of the detach NAS controller operation. This operation might take up to 30 minutes to complete.

After you detach a NAS controller:

- An X icon overlays the controller name in the PS Series group hierarchy.
- The NAS controller resets to its factory defaults and powers off, if possible.

NOTE: Wait for the detach NAS controller operation to complete before starting another long-running operation, such as an add NAS controller pair, attach NAS controller, or NAS reserve resize operation.

You can perform some member maintenance operations with no impact on member operation (for example, replacing a power supply). You can perform other NAS cluster maintenance operations by cleanly shutting down one NAS controller (resulting in a degraded pair), performing the maintenance, and then restarting the controller. See your NAS appliance hardware documentation.

If the detach NAS controller operation fails, contact your customer support representative.

Attach a NAS Controller

When a NAS controller is attached, it inherits the NAS cluster configuration settings of the original NAS controller.

Prerequisites

Before attaching a NAS controller to a controller pair, refer to your hardware documentation for information to obtain, and verify the following information and NAS controller physical state:

- Obtain the service tag identification number for the NAS controller that you want to attach.
- Power on the NAS controller and place it in standby mode.
- Connect all network ports to the appropriate switches, especially the SAN network ports.
- Enable IPv6 in your switch stack.
- Detach the old NAS controller before you attach the new one. You cannot attach a new controller if the old controller is down but still attached.

Procedure

To attach a NAS controller:

1. Click **Group Configuration**, then expand **Members**.
2. Select a NAS member and, if necessary, expand it to display and select the required NAS controller name.
3. Click **Attach NAS controller**.

4. In the **Attach NAS controller** dialog box, select the NAS controller that you want to attach to the NAS cluster. You can identify a NAS controller by its service tag.

   NOTE: If the NAS controller is not listed, verify the physical state and network connections for the controller and then click Rediscover to refresh the list of controllers in the dialog box.

5. Click the **Attach NAS controller** button.

A progress window opens, showing the progress of the attach NAS controller operation. This operation might take up to 30 minutes to complete.

NOTE:

- **Wait for the attach NAS controller operation to complete before starting another long-running operation, such as an add NAS controller pair, detach NAS controller, or NAS reserve resize operation.**
- **During an attach member or add member operation, or after a member restarts, the following warning might appear: `File system checker failed`. You can ignore this message.**

After attaching a NAS controller, you must rebalance the client connections.

If the attach NAS controller operation fails:

1. Verify that the NAS controller is correctly connected to the network and check each controller's physical status and network connections.
2. Verify that you are the only PS Series group administrator trying to configure the NAS controller into a NAS cluster.
3. Click **Retry**. If the attach NAS controller operation fails again, contact your customer support representative.

About Updating NAS Controller Firmware

The NAS firmware is supplied in service packs that you download onto the NAS controllers. You can update the NAS firmware in the Group Manager CLI. In some cases, you must update the firmware on the NAS controllers in a NAS cluster.

NOTE: Dell Storage Update Manager provides guided update management to simplify the process of updating the firmware for your EqualLogic platforms, including PS Series, FS Series, and disk drives. This tool facilitates both single-member and multimember group updates.

Update NAS Controller Firmware

Dell recommends that you perform this task during a maintenance window, because it can take 30 to 45 minutes for the update to complete. During this procedure, NFS will pause intermittently, but I/O will resume with no manual intervention.

Prerequisites

- You cannot update NAS controller firmware when a NAS controller is down or detached.
- You must use the grpadmin account password to update NAS controller firmware.
- All SMB clients will disconnect and reconnect when the NAS controller restarts. For this reason, Dell suggests that you stop all I/O to the NAS cluster.
- NFS pauses intermittently, but I/O will resume with no manual intervention.

Procedure

To update the NAS controller:

1. Download the service pack from the customer support website: eqlsupport.dell.com
2. For model FS7600-series NAS appliances, follow the instructions to copy the firmware on to a CD-ROM or a USB memory stick. If you use a CD-ROM, you need a USB CD-ROM reader connected to the NAS appliance.
3. Use FTP (enabled by default) to copy the service pack file to the NAS controller, specifying grpadmin credentials.

   Important: When using an FTP client, make sure that the file is transferred in binary mode. If the service pack is transferred using "auto" mode, the file is treated as text and transferred in ASCII mode. Transferring the file in ASCII mode adds control characters to the service pack, which might cause the embedded checksum to fail.

Upload the service pack by opening a URL using Windows Explorer (not Internet Explorer) or any other FTP client utility. For example:

```
ftp://grpadmin@nas_cluster_management_ip_address:44421/service_pack
```

Do not alter the service pack file name in any way.

4. When prompted, type the password for the grpadmin account.
5. From a server with access to the client network, use ssh to connect to the NAS cluster management IP address (not the PS Series group management address, if configured) and log in to the NAS cluster. For example:

   ```
   ssh grpadmin@nas_cluster_management_IP_address
   ```

   When prompted, type the password for the grpadmin account.
6. At the CLI prompt, enter the following command, specifying the name of the service pack file that you copied to the NAS controller. One at a time, each NAS controller will update and reboot.

   ```
   service-pack start service_pack_file_name
   ```

   The current ssh session terminates when the first updated NAS controller reboots.

   NOTE: You cannot reestablish an ssh session to the NAS cluster until the firmware update completes on all the NAS controllers.
7. Perform a mass rebalance operation to rebalance the SMB clients between NAS controllers:
   a. In the GUI, click **Group**, expand **Group Configuration**, and then select the NAS cluster.
   b. Click the **Advanced** tab.
   c. In the SMB Client Connections panel, click **Rebalance connections**.
8. Resume the NAS cluster. All management and I/O operations can continue.
9. After all the NAS controllers reboot and the firmware update completes on all the NAS controllers, from a server with access to the SAN network, use ssh to connect to the NAS cluster IP address and log in to the NAS cluster (as described in step 5).
10. Enter the following command at the CLI prompt to ensure that the update operation succeeded:

    ```
    service-pack status
    ```
11. Enter the exit command to log out of the NAS cluster.

If the firmware update operation fails, contact your customer support representative.

Cleanly Shut Down a NAS Controller Pair

For some NAS controller maintenance tasks, you might need to cleanly shut down and power off a NAS controller pair.

To cleanly shut down a NAS controller pair:

1. In the Group Manager GUI, identify the service tag identification numbers for the controller pair.
2. Refer to your hardware documentation to determine how to find the service tag information for a specific NAS appliance. This number might be in the controller firmware (software accessible) or on a physical label or tag on the NAS appliance chassis.

The following manuals address service tag location and how to shut down a NAS controller pair:

- *Dell Equallogic FS7500 Series Appliances Hardware Owner's Manual*
- *Dell Equallogic FS7600 Series Appliances Hardware Owner's Manual*
- *Dell Equallogic FS7610 Series Appliances Hardware Owner's Manual*

14

NAS Container Operations

Table 44. Basic and Advanced NAS Container Operations provides a list of basic and advanced NAS container operations.

Table 44. Basic and Advanced NAS Container Operations

Basic	Create, modify, or delete NAS containers Create, modify, or delete an SMB share Enable, show, modify, or disable an SMB home share Create, modify, or delete an NFS export
Advanced	Rebalance SMB client connections across NAS containers Add, modify, or delete a NAS antivirus server Protect NAS container data with NDMP Modify the file security style for NAS containers Create, modify, or delete quotas Enable thin provisioning and data reduction Reinstall the FluidFS operating system

Create a NAS Container

1. Click **NAS**.
2. In the Activities panel, click **Create NAS container**. The Create NAS Container wizard starts.

NOTE: The NAS replication container name field accepts names that are up to 229 bytes in length. However, if a source container name is longer than 225 bytes, the last few bytes will be deleted so that the system can add a 4-byte numeric identifier.

If the source container name includes Unicode characters and is longer than 225 bytes, the entire name string is replaced with the string `dest_` plus a 4-byte numeric identifier.

Ensure that no NAS volumes have the following names. These names are reserved for internal FluidFS cluster functions.

- ..
- .snapshots
- acl_stream
- cifs (or smb)
- int_mnt
- unified
- Any name starting with locker_

Modify NAS Clusterwide Default NAS Container Settings

To modify or display the NAS clusterwide default NAS container space settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default NAS Container Settings** panel, modify the settings as needed.

Modify NAS Clusterwide Default NAS Container Permissions

To modify or display the NAS clusterwide default NAS container permission settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default NAS Container Permissions** panel, modify the settings as needed.

NOTE: If the selected style is UNIX, default UNIX permissions are shown for files and directories that were created in the NAS container by Windows clients using SMB.

Modify NAS Clusterwide Default NFS Export Settings

To modify or display the NAS clusterwide default NFS export permission and trusted user settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default NFS Export Settings** panel, modify the settings as needed.

Modify NAS Clusterwide Default SMB Share Settings

To modify or display the NAS clusterwide default SMB share settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default SMB Share Settings** panel, modify the settings as needed.

Modify a NAS Container Name

1. Click **NAS**, expand **NAS clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Modify settings**.
3. In the Modify Settings dialog box, click the **Container Name** tab.
4. In the **Name** field, type the new name for the NAS container.

   The NAS container name must be unique and can contain up to 230 characters, including letters, numbers, and underscores. The first character must be a letter or an underscore.

   Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. Click **OK**.

Modify the Size of a NAS Container

When you increase the NAS container size, the maximum size is the amount of available free space in the NAS reserve.

To increase or decrease the size of a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.

2. Click **Modify Settings**. The Modify Settings dialog box opens.
3. Click the **Space** tab.
4. In the **Size** field, enter the new size for the NAS container.
5. Click **OK**.

Modify the Snapshot Reserve and Warning Limit for a NAS Container

NOTE: Select the Enable data reduction checkbox to activate the Modify policy button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.

To modify the snapshot reserve or the snapshot reserve in-use space warning limit for a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Modify snapshot settings**.
3. In the **Modify snapshot settings** dialog box, you can:
   - Specify the new snapshot reserve value, as a percentage of the NAS container reserve. The maximum snapshot reserve is 90 percent of the NAS container reserve.
   - Use the slider to adjust the snapshot reserve warning limit. The NAS container generates an event when the used snapshot reserve reaches the specified percentage shown under the slider. The minimum in-use space warning limit is 10 percent of the snapshot reserve; the maximum is 100 percent.

     NOTE: NAS Container Snapshot Warning Threshold Display Error—For NAS containers with snapshot reserve warning thresholds below 10%, the warning threshold value is not properly displayed. This issue is caused by a rounding error in an internal process, and does not apply to warning thresholds set above 10%.
4. Click **OK**.

Modify the In-Use Space Warning Limit for a NAS Container

The NAS container generates an event when in-use NAS container space reaches the percentage specified by the warning limit.

To modify the NAS container in-use space warning limit:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Modify settings**.
3. In the Modify Settings dialog box, click the **Space** tab.
4. In the **NAS container size** bar graph, adjust the slider to specify the new value for the in-use space warning limit.
5. Click **OK**.

Modify a NAS Container for Few Writers Workloads

The few writers workloads feature defines the way files (inodes) are distributed between domains to optimize usage on the NAS cluster. This feature can be enabled or disabled, on a per NAS container level. The default setting is disabled.

To enable a NAS container for few writers workloads:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Modify settings**.
3. In the Modify Settings dialog box, click the **Advanced** tab.
4. Select the **Enable optimization** checkbox.
5. Click **OK**.

Delete a NAS Container

If you no longer want the client data in a NAS container, you can delete the NAS container. After deleting a NAS container, the NAS container space becomes free NAS reserve space.

NOTE: Deleting a NAS container deletes all the snapshots of the NAS container and all the SMB shares and NFS exports in the NAS container.

After they are deleted, recovery containers cannot be recovered using the volume recovery bin.

When you delete a NAS container, its replica is not deleted; the replica is promoted to a container on the destination cluster.

To delete a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Delete container**.

   NOTE: You must disable the SMB home shares before you can delete a NAS container that has SMB home shares defined in it.
3. Click **Destroy *size* data**, where *size* is the size of the NAS container that you want to delete. You must select this checkbox to delete the container.
4. Click **Yes**.

NFS Netgroups

An NFS (Network File System) export provides an effective way of sharing files and data across networks. FluidFS provides support for the SMB and NFS protocols. NFS is the network file-sharing protocol that is associated with the Linux and UNIX operating systems. An NFS export can be accessed by using NFS v3 or NFS v4. When specifying restrictions for NFS, you can apply the following settings:

- A single IP address
- An IP network
- A netgroup
- No restrictions

You can configure netgroups and set restrictions so that only specific netgroups can access an NFS export. Of the three fields in the netgroup form (host, user, domain), FluidFS supports only the host field.

You can configure the netgroups host field with one of the following items:

- A single IP address
- A host name

NOTE: Group Manager supports host netgroup (host name or IP address) only. User netgroup, domain netgroup, and subnet in host netgroup are not supported at this time.

Network file sharing works with access protocols in the following ways:

- Homogeneous – All file sharing is done through a single protocol, as in a Windows-only file-sharing environment.
- Heterogeneous – Both SMB and NFS clients access files and directories.

NOTE:

- **To use the NFS protocol, UDP traffic on port 2049 must be allowed between a client and the NAS cluster.**
- **Any files or directories stored in a FluidFS NAS volume includes metadata that saves access permissions to the file or directory.**

Access NFS Exports

NFS version 3 relies on client authentication services. If a user can authenticate to a client machine, they can use the NFS mount command to access a FluidFS export. During an NFS export session, each operation is verified using the user's UID (user ID) and GID (group ID). FluidFS queries the configured directory services to obtain the group membership of the UID provided. NFS version 4 uses Kerberos authentication services.

Initially, only a root user can access an NFS export. To set the ownership and permissions for the share:

1. Ensure that the NFS export has read-write permissions. Also, make sure that the trusted user setting is `All`.
2. For security, type the IP address of the export client in the **Limit access to IP address** field. This action ensures that only the client's root user can access the export.
3. From a Linux or UNIX client, enter the showmount command to display the NFS exports that are hosted by the NAS cluster. For example: `showmount -e nas_vip`
4. As a root user, mount the NFS export and specify the NAS service IP address, the file-system name, and the export path that you obtained from using the showmount command. Then, designate the local mount point. For example:

   ```
   mount [options] NAS_service_ip_address:/file_system_name /export_path mount_path
   ```

   Dell recommends the following options:

   ```
   hard,tcp,nfsvers=3,timeo=3,retrans=10,rsize=32768,wsize=32768
   ```
5. For the NFS export, you can enter the appropriate owner, group, and permissions. Use the chown, chgrp, and chmod commands, respectively.

You can use NFS and NFS exports in the following ways:

- To use NFS over UDP, adjust the firewall to allow a source IP from any of the NAS controllers, then open the firewall to allow for port ranges.
- To prevent root from writing to the NFS export, modify the export and change the trusted user setting to `All except root`.
- To make the NFS export read-only, modify the export and change the permission to `Read-only`.
- The NFS protocol is case-sensitive. For example, two files with the same name can reside in the same folder, even if their names differ only by a single uppercase or lowercase character.
- Any file or directory stored in a FluidFS NAS volume includes metadata that saves access permissions to the file or directory.

Create an NFS Export

To create an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the NAS panel, right-click the container that you want to create the export in and select **Create NFS Export**. The Create NFS Export wizard opens.
3. In the General Settings dialog box, specify the **NFS export name, directory**, and select the appropriate **NFS file id compatibility** setting (32-bit or 64-bit). Click **Next**.

   NOTE: Starting with NFS v3, all file IDs are 64-bit. However, some older applications cannot properly handle file IDs where the upper 32 bits are utilized. To solve this problem, you can use either 32-bit or 64-bit file ID compatibility. The file ID remains 64-bit, but only the lower 32 bits are utilized.

   32-bit file ID compatibility provides maximum client compatibility. It works with both legacy applications and new applications. 64-bit file ID compatibility is used for clients where 64-bit file IDs are acceptable and where the set of files that are available in the system is very large.
4. In the Create NFS Export – Access Permissions dialog box, select a Client Access option to specify which client machines (**All Clients, Limit Access to IP** , or **All clients in a netgroup**) are allowed to access the NFS export. Click **Next**.
5. Specify whether the Permission level is **Read-write** or **Read-only** for the export.
6. Select the Trusted users setting that is appropriate for your configuration (options are **All, All except root**, and **Nobody**). Click **Next**.

   NOTE: When you create a new NFS export, select All in the Trusted users area to ensure that you have access until the configuration is complete.
7. The summary of the new NFS export is displayed. Use the Copy command to copy and save your settings for future reference. Click **Finish** to save your selections. Click **Back** to make changes.

Modify the Client Access Setting for an NFS Export

To modify the client access setting for an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.

3. Select the NFS export in the NFS Exports panel and then click **Modify NFS Export**. The Modify NFS Export dialog box opens.
4. In the dialog box, click the **Permissions** tab.
5. Specify whether to allow access to all clients, if they meet other access control requirements, or to a specific IP address or subnet. You can use asterisks in the IP address.
6. Click **OK**.

Modify the Permission for an NFS Export

To modify the permission (read-write or read-only) for an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.
3. Select the NFS export in the NFS Exports panel and then click **Modify NFS Export**. The Modify NFS Export dialog box opens.
4. In the dialog box, click the **Permissions** tab.
5. Specify the permission.
6. Click **OK**.

Modify the Trusted Users for an NFS Export

To modify the trusted users for an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.
3. Select the NFS export in the NFS Exports panel and then click **Modify NFS Export**. The Modify NFS Export dialog box opens.
4. In the dialog box, click the tab.
5. Specify the trusted user setting as either:

   - all except root (sometimes referred to as “root squash”)
   - nobody (or none)
   - all

   NOTE: Before setting the trusted user to all, you must set the Limit access to IP addresses field. Otherwise, the all selection appears dimmed.
6. Click **OK**.

Modify NAS Clusterwide Default NFS Export Settings

To modify or display the NAS clusterwide default NFS export permission and trusted user settings:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Defaults** tab.
3. In the **Default NFS Export Settings** panel, modify the settings as needed.

Modify an NFS Export Directory

If you modify the directory of an NFS export, clients that have the NFS export mounted will not be able to read from or write to the export and will see `stale NFS handle` messages. Clients must unmount the export and then mount it using the new directory.

To modify an NFS export directory:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.
3. Select the NFS export in the NFS Exports panel and then click **Modify NFS export**. The Modify NFS Export dialog box opens.
4. In the dialog box, click the tab.
5. In the **Exported directory** field, specify the new NFS export directory.

6. Click **OK**.

Modify an NFS Export

NOTE: To edit an NFS export, you must have group administrator (grpadmin) privileges.

To edit the properties of an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.
3. Select the NFS export in the NFS Exports panel and then click **Modify NFS export**. The Modify NFS Export dialog box opens.
4. In **General settings**, you can edit only the NFS file ID compatibility property and directory name. You cannot change the name of an NFS export.

   NOTE: The OK button in this dialog box appears dimmed until you enter data in the mandatory fields identified by an asterisk (*).
5. In **Access Permissions**, select appropriate options to edit:
   - Clients that must access the NFS export
   - Access-level permissions
   - IP address range to which the user rights must be assigned
6. To change permission-related properties, click the **Permissions** tab.
7. Click **OK**.

   The NFS export properties are updated and displayed under the **NFS Exports** tab.

   NOTE: If you create an NFS export as part of creating a NAS container, you must modify the export to change the access and security settings.

About NFS Export Security Methods

FluidFS provides optional security methods for a NAS cluster on a per NFS export basis. Use the CLI to specify these methods. See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information.

The nas-cluster select container select nfs-export select security-methods command allows you to choose from four security methods. Values for the `security-methods` parameter are:

- `sys` – UNIX style
- `krb5` – Kerberos v5
- `krb5i` – Kerberos v5 integrity
- `krb5p` – Kerberos v5 privacy

NOTE: By default, all security methods are enabled. If you specify one or more parameter values, you will override the default setting, thus disabling the remaining unspecified values. For example, if you create only `sys` and `krb5p` security methods for the specified cluster, you disable `krb5` and `krb5i`.

Delete an NFS Export

To delete an NFS export:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **NFS Exports** tab.
3. Select the NFS export name in the NFS Exports panel.
4. Click **Delete NFS export**.
5. Confirm that you want to delete the NFS export.

About SMB Shares

SMB shares provide an effective way to share files located on a FluidFS cluster, such as the FS76x0, by using the Server Message Block (SMB) protocol. FluidFS v4 supports SMB protocol versions 1.0, 2.0, and 3.0. The default SMB protocol version is SMB 3.0. You can set the default to an earlier version using the CLI command **nas-cluster select** *cluster_name* **smb-protocol**. Refer to the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about this command.

An SMB *share* is an access point to files and folders stored on a NAS container. An SMB share points to a folder on a NAS container (/NASFolder) or to the root of the NAS container (/...). Keep the following considerations in mind about SMB shares:

- You can create a share for one user or for multiple users.
- You determine a user's access to shares by specifying permissions in the Windows Explorer security settings.
- You can configure SMB shares in the Group Manager GUI or CLI.

Access SMB Shares in Windows

To access an SMB share, you must have a valid user name and password. Authentication can be either local or remote.

Before you can write to the SMB share:

- The group administrator must set the SMB administrator password.
- An administrator must log in to the SMB administrator account on the share and assign you write permission to the SMB share through the Windows operating system.

  If the NAS cluster is part of an Active Directory domain, you can perform the operation through the domain administrator account.

Access Shares

NOTE: Depending on your Windows version, these steps might vary slightly.

1. Click **Start** → **Run**.
2. In the **Open** field, specify the NAS cluster IP address (or the DNS name associated with the NAS cluster IP address) and click **OK**.
3. Right-click the share and select **Map Network Drive**.
4. In the **Map network drive** dialog box:
   a. Enter `\\`*`service_ip_address`*`\`*`share_name`* or `\\`*`service_dns_name`*`\`*`share_name`*.
   b. Click **Connect using a different user name**.
5. Click **Finish**.
6. In the **Connect as** dialog box, enter a valid user name and password and click **OK**.

You can enter `SMBstorage\administrator` for a user name and the associated SMB administrator password. The SMB administrator has write permission on all SMB shares by default.

If you have joined the NAS cluster to an Active Directory domain, you can also enter *`domain_name`*`\administrator` for a user name and the domain password. The domain administrator has write permission on all SMB shares.

Mount a NAS SMB Share from UNIX

To mount an SMB share from a UNIX operating system, use one of the following commands:

```
smbmount //service_ip_address/share_name /local_directory -o user_name=user_name
```

```
mount -t smbfs -o user_name=user_name,password=password //service_ip_address/share_name /
local_directory
```

Create an SMB Share

To create an SMB share:

NOTE: You can create an initial SMB share when you create a NAS container. However, you cannot configure and enable the NAS antivirus service. You must modify this initial SMB share to configure and enable the antivirus service.

1. Click **NAS**, expand **NAS clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Create SMB share** to open the wizard.
3. In the **General Settings** page:

   a. Type a name for the SMB share in the **Name** field. An SMB share name can contain up to 24 characters, including letters, numbers, dollar sign ($), and underscores. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

      NOTE: Do not create a regular SMB share with the name `homes`. If you attempt to create a share with this name, the following error is displayed: `Share name is reserved for home shares.`

   b. Type the full path of the shared directory in the **Directory** field. Directory names can contain 256 characters, including one preceding slash. The full path for an SMB share can contain up to 511 characters, including additional slashes.

      Click **Next**. If the directory does not exist, it will be created automatically. You cannot modify the SMB directory.
4. Select the **Enable Access-based Enumeration** checkbox to enable this feature. This feature allows users of Windows Server 2003–based file servers to list only the files and folders to which they have access when browsing content on the file server. This restriction eliminates confusion that can occur when users connect to a file server and encounter a large number of files and folders that they cannot access.
5. In the **Antivirus Settings** page, virus scanning is enabled by default. Optionally, you can also specify:

   - File extensions to exclude
   - Directory paths to exclude
   - Size of files to exclude from virus scanning
   - User access to large unscanned files

     (When an SMB file is accessed by a client and it needs to be scanned, it is sent over the network to an antivirus scanner. In some cases, the antivirus server has limits on the size of the files that it can handle or, in other cases, an administrator might want to exclude frequently modified files from scanning for performance reasons.)

   Select the options that you want and click **Next**.
6. In the **Summary** page, you can review your settings. To make changes, click **Back** to return to the previous screens. If your settings are correct, click **Finish** to create your SMB share.

Set the SMB Password

Clients use the SMB password to access SMB shares.

To allow client access to SMB shares, you must set the SMB administrator password. If you do not set this password, you cannot properly set the ownership and permissions for the SMB shares.

NOTE: The local administrator built-in account (Administrator), which appears in the list of local users in the NAS Cluster – Local Users and Groups window, is the same account as the SMB administrator.

To set or modify the SMB password:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Authentication** tab.
3. Click **Modify SMB administrator password**.
4. In the **SMB Administrator Password** dialog box, specify the password and confirm it.
5. Click **OK**.

Modify an SMB Share Directory

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name..

2. Click the **SMB Shares** tab.
3. Select the SMB share in the SMB Shares panel and click **Modify SMB share**.
4. In the Modify SMB Share dialog box, click the **General** tab.
5. Modify the directory as needed.
6. Click **OK**.

Delete an SMB Share

To delete an SMB share and delete all the user data stored in that share:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. Select the share name in the SMB Shares panel and click **Delete SMB share**.
4. Confirm that you want to delete the SMB share.

Rebalance SMB Client Connections Across NAS Controllers

To rebalance SMB client connections across all the available NAS controllers:

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Advanced** tab.
3. In the SMB Client Connections panel, click **Rebalance connections**.
4. Confirm that you want to rebalance client connections.

Enable or Disable SMB Message Signing

To help prevent attacks that modify SMB packets in transit, the SMB protocol supports the digital signing of SMB packets. SMB message signing can be enabled or disabled using the nas-cluster select *cluster_name* smb-protocol message-signing command in the Group Manager CLI. For more information about the CLI commands, see the *Dell EqualLogic Group Manager CLI Reference Guide*, which is available on the support site.

Enable or Disable SMB Message Encryption

SMB3 adds the capability to make data transfers secure by encrypting data in-flight, to protect against tampering and eavesdropping attacks. SMB message encryption can be enabled or disabled using the nas-cluster select *cluster_name* smb-protocol message-encryption command in the Group Manager CLI. For more information about the CLI commands, see the *Dell EqualLogic Group Manager CLI Reference Guide*, which is available on the support site.

Modify SMB Share NAS Antivirus Settings

A newly created or existing SMB share inherits the NAS clusterwide default antivirus settings. You can modify these settings on individual SMB shares as needed. For example, you might want a specific Threat Action policy or you might want to exclude certain directories from scans to improve performance.

If you unconfigure the NAS antivirus service, you cannot modify the antivirus settings for SMB shares. However, if you subsequently reconfigure the service, the settings that were previously in place are preserved.

To create specific NAS antivirus settings for a particular SMB share:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. Select the SMB share in the **SMB Shares** panel.
4. Click **Modify SMB share** to open the dialog box.
5. Click the **Antivirus Settings** tab.
6. Modify the NAS antivirus settings as follows:

- Enable or disable virus scanning.
- Modify file extensions.
- Select directory paths to exclude.

7. Click **OK**.

NOTE: The default antivirus exclude path is no longer available. Directory paths must already exist in the SMB share before they can be excluded.

To exclude directory paths from antivirus scanning:

1. Create the SMB share without the exclude option. See Create an SMB Share.
2. Go to the SMB share and create the directory paths that you want to exclude from antivirus scanning.
3. In the Group Manager GUI, modify the antivirus setting by specifying the file extensions and/or the directory paths to exclude.

CAUTION: Any default antivirus exclude path that you have set in previous versions will be lost.

Access-Based Enumeration

Access-based enumeration provides ease of use to users and reduces security risks. Access control is managed by using an access control list (ACL). An ACL is a set of rules that defines the access rights for a file folder in an SMB share.

NOTE: To edit ACL rules on files and folders that are located on the FluidFS cluster, you must use Microsoft Management Console 3.0 (MMC 3.0) or later.

By using access-based enumeration, you can access files and folders in a share that have sufficient permission on the basis of file and folder ACLs. Users can view only those files for which they have enough permissions to access (for example, read-only access). Users will not be able to see the files to which access has not been configured.

To enable or disable access-based enumeration, you must have grpadmin privileges. You can also enable access-based enumeration from the Group Manager CLI. For more information about the CLI commands, see the *Dell EqualLogic Group Manager CLI Reference Guide*.

Enable Access-Based Enumeration at the SMB Share Level

To enable access-based enumeration at the individual SMB share level:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click **Create SMB share** in the Activities panel.
3. In the Create SMB Share wizard, select the **Enable Access-based Enumeration** checkbox.
4. Click the Save all changes icon.

Enable Access-Based Enumeration as a Default Property on SMB Shares

NOTE:

- **Access-based enumeration can be enabled or disabled at the SMB share level.**
- **To enable or disable access-based enumeration as a default property at a container level, you must have grpadmin privileges.**

To enable access-based enumeration on all the newly created SMB shares of a NAS cluster:

1. Click **Group** → **Group Configuration** and then select the group to configure access-based enumeration.
2. Click the **Defaults** tab.
3. In the Default SMB Share Settings panel, select the **Enable Access-based Enumeration** checkbox.
4. Click the Save all changes icon.

NOTE: Only the SMB shares created in a NAS container after setting this default property will have access-based enumeration by default. SMB shares that were created before setting this property still have the properties that were set when the shares were created.

Enable Access-Based Enumeration on Newly Created SMB Shares

To write data to a NAS container, you must create an SMB share.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Create SMB share** to open the dialog box.
3. Type a name for the SMB share and directory and then select the **Enable Access-based Enumeration** checkbox.
4. Click the Save all changes icon.
5. Click **Next** to display the Antivirus Settings page.
6. Type or select appropriate data in the fields, and then click **Next** to display the Summary tab.

   The Antivirus Settings tab opens only if you have configured an antivirus server to the NAS cluster. In the Antivirus Settings section, if Status displays `not configured`, you probably have not configured an antivirus server for your NAS cluster. For information about configuring an antivirus server, see About NAS Antivirus Servers.
7. The Summary tab displays information about the SMB share that will be created. Review the information and click **Finish**. To make any changes, click **Back**.

   NOTE: An SMB share folder is created and displayed on the SMB Shares tab.

Enable Access-Based Enumeration as a Default Option for New SMB Shares

To enable access-based enumeration option for individual SMB share levels:

1. Click **Group** → **Group Configuration** and then select the group to configure access-based enumeration.
2. Click the **Defaults** tab.
3. Select the **Enable Access-based Enumeration** checkbox in the Defaults SMB shares settings panel.
4. Click the Save all changes icon.

Modify Access-Based Enumeration SMB Share Properties

NOTE: To edit the Access-Based Enumeration feature, you must have (grpadmin) privileges.

To edit the access-based enumeration SMB share properties:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. Double-click the SMB share for which you want to modify access-based enumeration. The **Modify SMB share** dialog box opens.
4. Select or clear the **Enable Access-based Enumeration** checkbox.

NOTE: You can enable access-based enumeration on the newly created SMB shares. For more information, see Enable Access-Based Enumeration on Newly Created SMB Shares.

Properties of the SMB share folder are updated and displayed in the panel on the **SMB Shares** tab.

About SMB Home Shares

When you create an SMB share, you create an actual share that points to a folder that is located on the NAS container. The SMB home share feature allows the administrator to create a single template that is used to automatically create a user home folder in the NAS cluster for each Windows user. Users can access their home folder by connecting to \\cluster. A virtual share is automatically created, pointing to the user home folder (which could be automatically created on the first login based on the home share setting).

NOTE: Automatic home folder creation cannot be enabled on an SMB home share in a NAS container with UNIX security style.

For example, when client **jsmith** connects to the FluidFS cluster, the feature will present **jsmith** with any available SMB shares, as well as an SMB home share labeled **jsmith**.

NOTE: You still must create the user folders yourself and set the permissions manually or by using an automated script if the automatic home folder creation option is not enabled in the SMB home share settings. See the following topics:

- Create Parent and User Folders
- Set User Permissions

SMB Home Shares Procedure Summary

This section provides a summary of how to enable SMB home shares and a link to each task.

Prerequisites

- Create a NAS container. See Create a NAS Container.
- If automatic home folder creation is disabled, create an SMB share that will be the root of all the users' home folders. The root will be the path-prefix value in the SMB home share setting. See Create an SMB Share.

Each of the three main tasks for creating an SMB home share has its own series of steps.

NOTE: Steps 1 and 3 are needed only when the automatic home folder creation option is disabled while enabling the SMB home share feature in step 2.

1. Give ownership of the SMB share (created previously) to the account that will create the folders for each user's home folder (for example, the domain administrator). See Give Ownership of the SMB Share to an Account.
2. Enable the SMB home share.
3. Connect to the SMB share that you created and create a home folder structure, under which you create a home folder for each user with appropriate permissions. The home folder structure must conform to the path prefix you chose. See the following topics:
    - Create Parent and User Folders
    - Set User Permissions

Give Ownership of the SMB Share to an Account

To give ownership to an account, follow these steps:

1. Select the Computer
2. Specify the Owner
3. Replace the Owner

Select the Computer

1. Open the **Shared Folders** MMC console.
2. Select **Connect to another computer** from the **Action** menu.

   The **Select Computer** dialog box opens.
3. In the dialog box, click the **Another computer** button and enter the NAS cluster name in the text box (for example, fluidfs.venus.local).
4. Click **OK**.

Specify the Owner

1. In the right pane of the **Shared Folders** window, double-click **Shares**. Then, right-click `home` (the SMB share that you created) to select **Properties** from the pop-up menu.

   Windows Explorer opens the **Properties** window.
2. On the **Share Permissions** tab, click the **Add** button.

   A selection dialog box opens for user, computer, service account, or group.
3. In the text box labeled **Enter the object name to select**, type a name (for example, `domain administrator`) and click **OK**.
4. Ensure that the object (domain administrator) is selected in the **Group or user names** box. The **Permissions** box displays the access rights for this object. Select the **Allow** checkbox to give **Full control** to the object.

Check marks appear in all the **Allow** checkboxes. This object is the share from which you will be creating a folder for each user's home share.

5. In the **Group or user names** box, select **Everyone**, then click **Remove** → **Apply**.

Replace the Owner

1. In the **Properties** window on the **Security** tab, click the **Advanced** button to open the window for **Advanced Security Settings for home**.
2. In the **Advanced Security Settings** window, click the **Owner** tab and then click the **Edit** button.
3. On the **Edit** window, click **Other users or group**.

   A selection dialog box opens for user, computer, service account, or group.
4. In the text box labeled **Enter the object name to select**, type a name (domain administrator) and click **OK** to return to the **Owner** tab.
5. Ensure that the checkbox is selected for **Replace owner on subcontainers and objects** and click **OK**.

   A message is displayed: `If you have just taken ownership of this object, you will need to close and reopen this object's properties before you can view or change permissions.`
6. Click **OK** in the message box.
7. On the **Owner** tab, ensure that the object that you entered is selected in the list box labeled **Change owner to** and click **OK**.
8. Click two **OK** buttons to close the windows and return to the Windows Explorer main window.

Create the SMB Home Share Using the CLI

In the CLI, use the following command to enable the SMB home share:

```
nas-cluster select cluster_name smb-home-share create container_name path_prefix
[parameters]
```

See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about the variables and parameters in this command.

NOTE: The default antivirus exclude path is no longer available. Directory paths must already exist in the SMB home share before they can be excluded.

To exclude directory paths from antivirus scanning:

1. Create the SMB home share without the exclude parameter.
2. Go to the SMB home share and create the directory paths that you want to exclude from antivirus scanning.
3. Modify the antivirus settings by specifying values for the following parameters:
   - `avexcludedirs extension_list`
   - `avextensions directory_list`
   - `avlargefilesize size`

Enable SMB Home Shares

Enable the SMB home shares feature to automatically create an SMB home share for each user that is accessing the NAS cluster.

1. Click **NAS**, expand **NAS Clusters**, and select **SMB Home Share**
2. In the Activities panel, click **Enable SMB home share** to open the wizard.
3. In the General Settings page, select the **Enable automatic home folder creation** checkbox.
4. Click **Next**.
5. The **Summary** page opens.
6. Click **Finish**.

Disable SMB Home Shares

To disable SMB home shares:

1. Click **NAS**, expand **NAS Clusters**, and select **SMB Home Share**
2. In the Activities panel, click **Disable SMB home share**.

3. Click **Yes** to confirm.

Modify SMB Home Share Settings

To modify SMB home share settings:

1. Click **NAS**, expand **NAS Clusters**, and select **SMB Home Share**.
2. In the Activities panel, click **Modify settings**.
3. In the SMB Home Share General dialog box, you can enable or disable automatic home folder creation, and enable or disable access-based enumeration. You can also modify antivirus settings if antivirus is enabled.
4. Click **OK**.

Create an SMB Home Folder Automatically

FluidFS v4 enables the creation of home folders automatically while creating SMB home shares. When enabled, the home folder is automatically created inside the container when the SMB home share is first accessed. Automatic home folder creation is not allowed for a NAS container with UNIX security style.

To automatically create an SMB home folder:

1. Click **NAS**, expand **NAS Clusters**, and select **SMB Home Share**
2. In the Activities panel, click **Modify settings**.
3. In the SMB Home Share General dialog box, select **Enable automatic home folder creation**.
4. Click **OK**

Connecting to the SMB Share and Creating Folders

To connect to the SMB share and create folders, follow these steps:

1. Create Parent and User Folders
2. Set User Permissions

Users can access their home/user folders through Windows Explorer (for example, \\mysystem\chrisb).

NOTE: If an SMB home share is disabled, the user folder and data in the container is not deleted. To access the data in the folder, recreate the SMB home share with the previous home share settings values (path prefix and folder template).

Create Parent and User Folders

1. In Windows Explorer, right-click the SMB share that you created and select **Open** from the pop-up menu.
2. Right-click the path prefix and select **Folder** from the New menu.
3. Type a name for the folder.

   This folder becomes the parent folder and the Windows domain name if the user folder template is domain-user.
4. Right-click the name of the parent folder and select **Folder** from the New menu to create a user folder.
5. Type the user name for the folder.

   You can create multiple user folders at this point.

Set User Permissions

1. Right-click the user folder and select **Properties**.

   The **Properties** window opens.
2. On the **Security** tab, click the **Edit** button.

   The **Permissions** window opens.
3. In the **Permissions** window, click the **Add** button.

   The selection dialog box opens for user, computer, service account, or group.
4. In the text box labeled **Enter the object name to select**, type the name of a user and click **OK**.

   The **Permissions** box displays the access rights for this object.
5. Select the **Allow** checkbox to give **Full control** to the object.

Check marks appear in all the **Allow** checkboxes.

6. Click **OK**.
7. Click two **OK** buttons.

   Both windows close, returning you to the Windows Explorer main window.
8. Close the MMC console.

Create a NAS Thin Clone

During the creation of a NAS thin clone, you cannot modify its size, minimum reserve percentage, or in-use warning limit percentage. These settings can be changed later.

1. In the **NAS** panel, select a local container that contains a snapshot.
2. Select the snapshot.
3. Click the **Create thin clone** link in the Activities panel to open the Create Thin Clone wizard.
4. Either use the default container name or modify the container name.
5. Click **Next**.

   A **NAS container size** horizontal bar is displayed, along with a minimum reserve slider. The percentages in the slider are not editable on the create page. The size is inherited from the base container and can be modified after the thin clone is created.
6. Move the slider to adjust the percentages.
7. Click **Next**.

   To create an NFS export and an SMB share, type the name and directory information in the dialog box.
8. Click **Next**.
9. Review the **Create NAS Thin Clone** summary.
10. Click **Finish**.

Client Networks

The client network is used to provide client access to the SMB shares and NFS exports. The client network is assigned one or more virtual IP addresses (VIPs) that allow clients to access the NAS cluster as a single entity. The client VIP also enables load balancing between NAS controllers, and ensures failover in the event of a NAS controller failure.

NOTE: Dell recommends that you connect all the client network interfaces and distribute the connections across stacked network switches for high availability.

NOTE: Dell recommends that you separate the client and SAN network to avoid the following problems:

- Congestion caused by client traffic and storage traffic passing over one network
- Issues related to static routes configuration

You can configure up to eight secondary or additional client networks on your NAS cluster. This configuration allows client traffic to be separated and secured. You can configure VLANS for each client network that is available in the NAS cluster.

If the client network is a routed network, you can modify the client network configuration and add more NAS cluster IP addresses for additional DNS load balancing across multiple client networks. The total number of NAS cluster IP addresses depends on the bonding mode and NAS appliance models.

A default gateway IP address, bonding mode, and MTU are the same for all the client networks. You cannot configure a unique default gateway IP address, bonding mode, and MTU for individual client networks.

If client access to the NAS cluster is not through a router (the network has a "flat" topology), you can define one client VIP per NAS controller. If clients access the FluidFS system through a router, define a client VIP for each client interface port per NAS controller.

Optimal Virtual IP Assignment

To optimize availability and performance, client connections are load balanced across the available NAS controllers. NAS controllers in a NAS cluster operate simultaneously. If one NAS controller fails, clients are automatically failed over to the remaining controllers. When failover occurs, some SMB clients reconnect automatically, while in other cases, an SMB application might fail and the user must restart it. NFS clients experience a temporary pause during failover, but client network traffic resumes automatically.

You must configure the client network virtual IPs for optimal load balancing. If you do not optimally configure the client network virtual IPs, the system prompts you to assign an optimal number of virtual IP addresses for the client network. Use the guidelines in the following table to set up an optimal number of VIPs.

Bonding Mode	Cluster Configuration	Optimal VIPs
ALB Mode	Cluster with 2 FS7x00 controllers	8
ALB Mode	Cluster with 4 FS7610 controllers	8
LACP Mode	Cluster with 4 FS7500 controllers	4

NOTE: You can autofill the VIPs for the client network and NAS controller IP addresses.

Viewing or Modifying a Client Network

To view or modify an existing client network in a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Network** tab.
3. In the Client Network panel, select the network and click **Modify**.

   The Modify Client Network dialog box opens.
4. In the dialog box, you can modify the following fields:
   - **Name**

     NOTE: You cannot modify the name of a primary client network.
   - **Netmask**
   - **Use VLAN tagging**
   - **VLAN ID (1 – 4094)**
   - **Virtual IP addresses**
   - **NAS controller IP addresses**
5. Modify the fields as needed and click **OK**.

   The client network details are updated and displayed under **Client Network** on the Network tab.

Deleting a Client Network

To delete an existing client network in a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Network** tab.
3. Under Client Network, select the network and click **Delete**.

   The **Delete client network** dialog box opens.

   NOTE: You cannot delete a primary client network.
4. To delete the selected client network, click **Yes**.

   The client network is deleted from the NAS cluster and removed from the Client Network panel.

To cancel, click **No**.

Modifying Client Network Properties

Default values for gateway IP address, bonding mode, and MTU are shared among all the client networks for a NAS cluster. Depending on the bonding mode selection, a message prompts you about the change in the number of virtual IP addresses for your client networks and the need to change the virtual IP address settings for each client network to ensure that load balancing is optimal. To modify the client network properties:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. In the Activities panel, click **Modify client properties**.

   You can modify the following fields:

   - Default gateway
   - MTU size
   - Bonding mode

     NOTE: When you change the bonding mode of a client network, the optimal number of VIPs is changed. To retain optimal configuration settings, you must manually modify the VIPs of each client network.
3. Click **OK**.

   The default client network configuration is updated for each client network.

About NAS Antivirus Servers

You cannot actively determine the status of antivirus servers from within Group Manager. Therefore, you cannot know whether antivirus servers are present and working. If no antivirus hosts are available, the following error message is logged to Group events:

```
No antivirus hosts are accessible. Virus scanning is not possible. SMB shares configured
with antivirus scan will not be accessible.
```

NAS antivirus allocates scanning operations to the antivirus servers to maximize the available scanning bandwidth. The fewer available antivirus servers, the more time required to scan files.

If you have only one antivirus server, you cannot delete that server until you first disable NAS antivirus on all SMB shares.

How NAS Antivirus Protects Data

When an SMB share user (or program) requests a file from the NAS cluster, NAS antivirus passes the file to an antivirus server for scanning. If the file is virus free, NAS antivirus permits user access. NAS antivirus does not scan that file again, providing it remains unmodified since the last check. The scan operation is transparent to the file's user, subject to the availability of an antivirus server.

If the antivirus server reports an infected file, the file is automatically quarantined. This action prevents the virus from contaminating other data files.

Users see no indication that a file is infected. Instead:

- A file deletion returns a system-specific `file not found` state for a missing file, depending on the user's client computer.
- An access denial might be interpreted as a file-permission problem.

At this point, only NAS cluster administrators can recover an uninfected version of the file, or access and process the infected file.

To gain access to an infected file, you must connect to the SMB share through another share on which the NAS antivirus service is disabled. Otherwise, NAS antivirus recognizes the file as infected, and denies access. You might also access the file through an NFS export, because NFS does not support NAS antivirus.

NAS antivirus must be enabled to actively scan clusters or SMB shares. You can enable antivirus scans for a cluster or an SMB share at any time.

NAS Antivirus Server Specifications

The following requirements apply for antivirus servers:

- FluidFS version 3.0 or later must be loaded on the cluster.
- The server must be accessible by the network. Dell recommends that the server be located on the same subnet as the NAS cluster.
- The server must run certified ICAP-enabled antivirus software.

Add a NAS Antivirus Server

Before you can add and enable the NAS antivirus service and configure scanning on NAS SMB shares, you must:

- Install the supported antivirus software on the external antivirus servers.
- Set up the antivirus software as a network-accessible service according to the user instructions for the antivirus software.

When you have installed the antivirus software, you can then add up to four antivirus servers, making these servers available to the NAS cluster.

If an antivirus server becomes unavailable, NAS antivirus posts an event to the group. For information about monitoring events, see Monitor Events.

Before you add NAS antivirus servers, you need the following information:

- The host name or IP address of up to four network-reachable servers on which you have previously installed a supported antivirus server service.

  Use the Group Manager CLI to ping the server's IP address for verification. If a server is unreachable, you cannot add it.
- A port number for the service, or you can use the default port number of 1334.

NOTE: You need an alternate port number only if you intend to run multiple instances of the antivirus software on the same server (perhaps for use by other clients). However, Dell recommends a single instance per server, dedicated to the NAS antivirus service.

To add a server:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab and go to the Antivirus Servers panel.
3. Click **Add**.
4. Type the name or IP address of a computer providing the antivirus service.
5. Type a port number or select **Use default port (1344)**.
6. Click **OK**.

You can now configure and enable NAS antivirus scanning, as NAS clusterwide defaults or on individual SMB shares.

Modify a NAS Antivirus Server

You must add at least one, and up to four, antivirus servers as described in Add a NAS Antivirus Server. You can then modify antivirus servers as required.

To modify a NAS antivirus server:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab.
3. Click **Configure Antivirus Servers** to go to the Configure Antivirus Servers panel.
4. Select one or more servers from the list and click **Modify**. The Modify Antivirus Server dialog box opens.
5. In the dialog box, change the server name or IP address.

6. In the **Port** field, type the port number or click **Use default port**.
7. Click **OK** to confirm your changes.

Delete a NAS Antivirus Server

You cannot delete the last server unless you first allow any in-progress operations to complete and you disable NAS antivirus on all SMB shares.

NOTE: Reducing the number of available antivirus servers might affect file-access performance.

To delete a NAS antivirus server:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab.
3. Click **Configure Antivirus Servers** to go to the Configure Antivirus Servers panel.
4. Select one or more servers from the list and click **Delete**. When prompted, confirm the deletion.

About NAS Antivirus Clusterwide Defaults

NAS antivirus clusterwide defaults apply automatically to all newly created SMB shares.

The defaults do not apply to any existing SMB shares for which you have manually configured share-specific NAS antivirus settings.

If you unconfigure the NAS antivirus service, you cannot modify the clusterwide defaults. However, if you subsequently reconfigure the service, the defaults that were previously in place are preserved. In addition, you must configure an antivirus server before you can specify clusterwide defaults.

To enable NAS clusterwide antivirus defaults:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab and go to the Antivirus Defaults for SMB Shares panel.
3. Select **Enable virus scanning**.
4. Click the Save all changes icon at the top of the NAS cluster window to save your changes.

   NOTE: You cannot navigate away from the NAS cluster window without confirming that you want to discard or save your changes.

When you have enabled antivirus defaults, you can configure the ability to control which file types or directories are scanned.

Enable the NAS Antivirus Service for a Cluster

You can enable the NAS antivirus service for the entire cluster. When you configure or modify the default NAS antivirus settings for a NAS cluster, the defaults apply to all newly created SMB shares.

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab and go to the Antivirus Defaults for SMB Shares panel.
3. Select or clear the **Enable virus scanning** checkbox.
4. Click the Save all changes icon on the top task bar to save the current selections.

After you enable the antivirus service, you can:

- Exclude file extensions from the scan
- Exclude directory paths from the scan

Exclude File Types on a Cluster

1. Prepare a list of file types that you want to exclude from scans:

   - Include variants of those file types, such as docx and doc to specify Microsoft Word documents.

- Use only numbers, letters, underscores, and dollar signs ($) in the file types.

2. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
3. Click the **Advanced** tab and go to the Antivirus Defaults for SMB Shares panel.
4. In either the File Extensions to Exclude or Directory Paths to Exclude subpanel, click **Add** to open the Add List Item dialog box.
5. Specify a file type such as xls or ppt. Do not include the period (.) that separates a file type from the file name.
6. Click **OK** to add the extension to the list of excluded files.
7. Repeat this process to add more file extensions.
8. Click the Save all changes icon at the top of the window.

NOTE: You also use this process to add directory paths for exclusion from antivirus scans. Make sure you click the Save All Changes icon at the top of the window to complete the process.

Enable the NAS Antivirus Service on an SMB Share

You can enable the NAS antivirus service for individual SMB shares. To find out whether antivirus scanning is enabled for a particular SMB share, see Monitor the NAS Antivirus Service.

NOTE: You cannot configure the antivirus service when you create an SMB share at the same time as you create a NAS container. You must configure the antivirus service when you modify an SMB share.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab to display a list of SMB shares.
3. Select an SMB share.
4. Click **Modify SMB Share**, then click the **Antivirus Settings** tab.
5. Select or clear **Enable virus scanning**.
6. Click **OK**.
   In the list of SMB shares, the column titled **Virus scanning** now reads `enabled` for this share.

After you enable the antivirus service, you can:

- Control which file types are excluded from being scanned
- Control which directories are excluded from being scanned

Exclude File Types on an SMB Share

1. Prepare a list of file types that you want to exclude from scans:
   - Include variants of those file types, such as docx and doc to specify Microsoft Word documents.
   - Use only numbers, letters, underscores, and dollar signs ($) for the file types.
2. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
3. Click the **Advanced** tab and go to the Antivirus Defaults for SMB Shares panel.
4. In either the File Extensions to Exclude or Directory Paths to Exclude subpanel, click **Add** to open the Add List Item dialog box.

   NOTE: Wildcard characters (*) and question marks (?) are not supported in antivirus exclude paths.
5. Specify a file type such as xls or ppt. Do not include the period (.) that separates a file type from the file name.
6. Click **OK** to add the file type.
7. Repeat steps 4 to 6 to add more file types to exclude from the scan.
8. Click the Save all changes icon at the top of the window.

You can also configure the specifications for large file handling from the Antivirus Defaults for SMB Shares panel:

1. In the **Exclude files larger than** field, type the file-size limitation for the antivirus scan.
2. Select a unit (MB, GB, or TB) from the drop-down menu.
3. Click the Save All Changes icon at the top of the screen.

Monitor the NAS Antivirus Service

If you have configured NAS antivirus, you can monitor which SMB shares are using the service.

1. Click **NAS** , expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. In the **Virus scanning** column, determine which shares have virus scanning enabled or disabled.

NAS Directory Paths and File Types Scan

You can control what directory paths and file types are scanned across a cluster or on SMB shares as follows:

- Directory paths to exclude — Specified as a path such as /tmp/logs (alternatively, folders and subfolders)
- File types to exclude — Specified by file type (file extension)

To exclude directory paths and file types from antivirus scans across clusters and on SMB shares:

- Specify directory paths to exclude
- Specify file types to exclude

You can specify file type and path strings using the formats specified in Table 45. NAS Antivirus File and Directory Path Scanning Specifications.

Table 45. NAS Antivirus File and Directory Path Scanning Specifications

Entity	Format
Directory path string	Forward slash (/) separated string of directories (folders), such as /user/programs/data/tmp
File type string	Alphanumeric string such as docx or blend1

Table 46. Characteristics of File Types and Directories describes permitted characteristics of file types and directory paths.

Table 46. Characteristics of File Types and Directories

Characteristic	File Types	Directory Paths
Default data	No default	No default
Default setting	Scan all files	Does not exclude any file types
Max characters	254 characters	254 ASCII characters per line item (path)
Wildcards	None	Do not use wildcard characters (*) and question marks (?).

Modify a Directory Path for a Cluster

You can specify the default NAS antivirus settings for a NAS cluster. These defaults apply to all newly created SMB shares, but the configuration of existing shares does not change.

To modify a directory path for a cluster:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab and go to the Antivirus Defaults for SMB Shares panel.
3. Either:
   - Click **Add**, enter a directory path, and click **OK**.
   - Select a directory path and then either:
     - Click **Modify**, edit the directory path as required, and click **OK**.

NOTE: Wildcard characters (*) and question marks (?) are not supported in antivirus exclude paths.

  - Click **Delete** and then click **Yes** to confirm.

4. Click **OK**.
5. Click the Save all changes icon.

Exclude Directory Paths for an SMB Share

You can exclude directory paths for a specific SMB share when you create an SMB share or when you subsequently modify the NAS antivirus settings for an SMB share.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. Select a share and click **Modify SMB Shares**.
4. Click the **General** tab to determine the root of the mapped share.
5. Click the **Antivirus Settings** tab.
6. In **Directory paths to exclude**, click **Add** to open the Add List Item dialog box.
7. Type a folder path that is based on the root of the mapped share.

   NOTE: In some cases, such as with .zip files and other archives, the antivirus server's settings override the NAS antivirus threat policies. Infected files are processed according to the antivirus server's configuration. Dell recommends coordinating the antivirus server's own policies with those of the NAS antivirus server to ensure that the system takes consistent action when it finds infected files.
8. Click **OK**.
9. Repeat the previous steps to add additional directory paths.
10. Click **OK** to close the Modify SMB Share wizard.

Modify File-Type Filtering Clusterwide

When you configure or modify the default NAS antivirus settings for a NAS cluster, the defaults apply to all newly created SMB shares.

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Advanced** tab.
3. In the **Antivirus Defaults for SMB Shares** panel, either:
   - Click **Add**, specify a file type (extension) to filter from antivirus scanning, and click **OK**.
   - In **File extensions to exclude**, either:
     - Select an extension, click **Modify**, change the extension, and click **OK**
     - Select an extension, click **Delete** and then click **Yes** to confirm.
4. Modify other file types as needed.
5. Click the Save all changes icon.

Modify File-Type Filtering for an SMB Share

You can modify the NAS antivirus settings for an individual SMB share.

NOTE: Wildcard characters (*) and question marks (?) are not supported in antivirus exclude paths.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **SMB Shares** tab.
3. In SMB Shares, right-click a share and select **Modify SMB Share**.
4. In the Modify SMB Share dialog box, click the **Antivirus Settings** tab.
5. In **File extensions to exclude**, either:
   - Click **Add**, specify a file type (extension) to filter from antivirus scanning, and click **OK**.

    - Select an extension, click **Modify**, change the extension, and click **OK**.
    - Select an extension, click **Delete**, and click **Yes** to confirm.
6. Repeat to add, modify, or delete additional file types.
7. Click **OK**.

Antivirus Policy

Depending on the antivirus policy, a file could be deleted immediately or made inaccessible to users and programs.

If the antivirus server or NAS cluster's default setting causes file deletion, you can only recover a previous (uninfected) file. For example, you can recover a deleted file from a NAS container snapshot.

The antivirus policies that result only in access denial are:

- Quarantine file (default)
  Quarantining a file involves the following steps:
  - The file is copied to a **.quarantine** folder at the volume root and is accessible only by administrators.
  - The file name is changed to include a string that specifies the date and identifies the file as infected.
  - File permissions are changed to permit access only by the root user.
- Prevent file access
  The file remains in place, but file permissions are changed to permit access only by an administrator.

Access Infected Files

Only administrators can access infected files. Infected files can be accessed through an NFS export. However, NAS antivirus does not support NFS exports.

To access infected files:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
   The infected file will be quarantined or in a denied access state.
2. Create an SMB share with a directory path to the directory that contains the infected file (such as the **.quarantine** directory).

   CAUTION: Do not enable NAS antivirus on this new SMB share.
3. Move the file to a location where you can either safely scan and disinfect the file according to the instructions provided by the antivirus software vendor or otherwise dispose of the file.

Create a NAS Container Quota

Quotas define how storage space on a NAS container is allocated among users and groups of users. You can create a quota for a specific user, for each user in a group, and for all the users in a group, and you can notify those users when they cross defined thresholds in their allocated space.

To create a quota:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Quotas** tab.
3. In the Quotas panel, click **Add**. The **Create quota** dialog box opens.
4. In the **Create quota** dialog box, select the type of quota:
   - Single user — Quota applies to a specific user.
   - Each user in the group — Quota applies to each user in a group.
   - Group total — Quota applies collectively to all the users in the group.

   Click **Next**.

5. Select the user type.
6. In the User field, you can enter a user name (or the beginning of a user name) and click the **Search** button.
7. Select the user and click **Next**.
8. In the **Create quota – Quota settings** dialog box, specify the following configuration settings:
   - Quota size and units (MB, GB, or TB)
   - In-use space warning limit, as a percentage of the quota size

   NOTE: Specifying zero for a quota size and warning limit disables the quota.
9. Click **Next**.
10. if the quota configuration is correct, click **Finish** in the Create quota – Summary dialog box. Click **Back** to make changes.

Modify a NAS Container Quota

NOTE: You can disable a NAS container quota by setting the quota size and the warning limit to zero.

To modify a NAS container quota:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Quotas** tab.
3. In the Quotas panel, select the quota target and click **Modify**. The Modify Quota dialog box opens.
4. In the dialog box, specify the following configuration settings:
   - Quota size and units (MB, GB, or TB)
   - In-use space warning limit, as a percentage of the quota size
5. Click **OK**.

Modify the Default NAS Container Group Quota

By default, the default group quota size and warning limit are set to zero, which disables the quota.

NOTE: You cannot delete the default group quota. You can set the default group quota size and the warning limit to zero to disable the quota.

To modify the default group quota:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Quotas** tab.
3. In the Quotas panel, select *defgroup* and click **Modify**. The Modify Quota dialog box opens.
4. In the dialog box, specify the following configuration settings:
   - Quota size and units (MB, GB, or TB)
   - In-use space warning limit, as a percentage of the quota size
5. Click **OK**.

Modify the Default NAS Container User Quota

By default, the default user quota is set to zero, which disables the quota.

NOTE: You cannot delete the default user quota. You can set the default user quota size and the warning limit to zero to disable the quota.

To modify the default user quota:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Quotas** tab.

3. In the Quotas panel, select *defuser* and click **Modify**. The Modify Quota dialog box opens.
4. In the dialog box, specify the following configuration settings:
   - Quota size and units (MB, GB, or TB)
   - In-use space warning limit, as a percentage of the quota size
5. Click **OK**.

Delete a NAS Container Quota

You cannot delete the default group quota or the default user quota. To disable a default quota, set the quota size and the warning limit to zero.

To delete a NAS container quota:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Quota** tab.
3. In the Quotas panel, select the quota target and click **Delete**.
4. Confirm that you want to delete the quota.

About Quota Directories

A **quota directory** is a directory that accounts for the size of files and directories under it by enforcing limits on them.

A quota directory provides a more granular approach for controlling space usage by a specific directory within a NAS container. You can mark a directory as a quota directory to limit the total amount of space that the directory uses within the NAS container. Quota directories work with user and group quotas, along with any capacity restrictions applied to the NAS container itself.

Configuring Quota Directories

NOTE:

- **You can create multiple quota directories, but nesting quota directories is not supported.**
- **Quota rules can be set on empty directories only. After you set a rule, it can be edited or deleted. If you delete the quota rule, the directory reverts back to normal directory behavior.**

You configure quota directories in the Group Manager CLI using the following commands (quota directories cannot be configured in the Group Manager GUI):

- nas-cluster select *cluster_name* container select *container_name* quota-dir create *directory_name soft_limit hard_limit*

  NOTE: You must specify the maximum size for a particular directory with the *hard_limit* variable. Optionally, you can use the *soft_limit* variable to specify a size that is less than the hard-limit size.
- nas-cluster select *cluster_name* container select *container_name* quota-dir show
- nas-cluster select *cluster_name* container select *container_name* quota-dir modify
- nas-cluster select *cluster_name* container select *container_name* quota-dir delete

The create command sets the quota rule for a directory.

See the *Dell EqualLogic Group Manager CLI Reference Guide* for more information about these commands.

Quotas and NAS Containers

Administrators use different quota mechanisms to regulate and monitor space usage within the NAS containers, such as user quotas, group quotas, and quota directories.

User quotas are applied differently, depending on whether the promoted replica container uses external authentication, local authentication, or a combination of the two.

NOTE: Dell does not recommend using a combination of local and external authentication where replication and quotas are applied.

External Authentication

External authentication is managed on a server whenever a user logs in to a container in the same group as the server. Using external authentication, a user can log in to different containers in the group using the same user name and password. This authentication is performed with Active Directory, LDAP, or NIS, for example.

When a user logs in to a container using external authentication, the user and group quotas defined for that user and any groups that the user belongs to are applied to that user while logged in to that container.

Local Authentication

Local authentication is managed directly on a container. With local authentication, the user name and password that a user enters to log in to a container apply only to that container.

When a user logs in to a container using local authentication, the user and group quotas defined for that user and any groups that user belongs to are not applied to that user. The quotas for the local user and the primary group for that local user are applied to that user, with the user quota taking precedence over the group quota.

About NAS Thin Provisioning

Thin provisioning enables you to efficiently allocate storage space, while still meeting application and user storage needs. With a thin-provisioned NAS container, space is created based on actual container data usage, enabling you to “over-provision” NAS cluster storage space (provision more space than what is physically available).

When NAS containers are thin-provisioned, NAS cluster storage space is consumed only when data is physically written to the NAS container, not when the NAS container is initially allocated. When data is written to a container, it initially fills or consumes reserved space. When reserved space is exhausted, it begins to consume unreserved space.

Thin provisioning offers the flexibility to provision NAS clusters to account for future increases in usage. However, because it is possible for the storage space used by the NAS container to exceed the storage space allocated to the NAS reserve, you should monitor available capacity to ensure that the NAS reserve always has sufficient free space available.

You can also specify a portion of the NAS container (reserved space) that is dedicated to the NAS container (no other container can take the space). The total reserved space of all NAS containers cannot exceed the available capacity of the NAS reserve. If a file is deleted from an unreserved portion of a thin-provisioned NAS container, the free space as seen in the NAS cluster increases. The freed-up capacity is also visible and available to clients in the SMB shares or NFS exports.

NAS Container Storage Space Terminology

The following list defines terminology used in Group Manager related to NAS container storage space.

Term	Description
Size	Maximum size of NAS container defined by the storage administrator.
Used Space	Storage space occupied by writes to the NAS container (user data and snapshots).
Reserved Space	A portion of a thin-provisioned NAS container that is dedicated to the NAS container (no other container can take the space). The amount of reserved space is specified by the storage administrator. Reserved space is used before unreserved space.
Unreserved Space	A portion of a thin-provisioned NAS container that is not reserved (other containers can take the space). The amount of unreserved space for a NAS container is the NAS container size minus the reserved space in the NAS container.

Term	Description
Available Space	Storage space that is physically available for the NAS container. The available space for a NAS container is the amount of unused NAS container space (reserved and unreserved), provided that the NAS reserve has free space.
Oversubscribed Space	A portion of a thin-provisioned NAS container that is not available and not in use by the NAS container. The overcommitted space for a NAS container is the amount of space allocated (not necessarily reserved) to thinly provisioned containers that exceeds the available NAS reserve free space. With thin provisioning enabled, storage space is consumed only when data is physically written to the NAS container, not when the NAS container is initially allocated. This difference means that more storage space can be allocated for the NAS container than has been allocated to the NAS container itself.
Snapshot space	Storage space occupied by snapshots of a NAS container.
Data Reduction Savings	Storage space reclaimed as a result of data reduction processing.

About NAS Containers

To provision NAS storage, you can create multiple NAS containers in a NAS cluster. In a NAS container, you can create multiple SMB shares and SMB home shares and NFS exports. Access to shares and exports is through one or more NAS cluster virtual IP addresses.

The number and size of the NAS containers in a NAS cluster depend on the storage needs of your NAS clients and applications. You can increase or decrease the size of a NAS container as needed without disruption to the NAS clients accessing the NAS container.

In addition, NAS containers:

- Have robust security mechanisms
- Support user and group quotas
- Support snapshots for data protection
- Support thin clones
- Support NDMP for remote backups
- Support replication to remote FS Series NAS clusters for disaster tolerance
- Support thin provisioning

You can create a single, large NAS container, or you can create many NAS containers. Creating multiple NAS containers enables you to apply different management policies to the NAS containers, as required by your organization. For example, you can apply different backup, snapshot, security, and quota policies to the NAS containers.

When you create a NAS container, SMB share, SMB home share, or NFS export, the NAS cluster applies default values.

Enable Thin Provisioning on a Container

To enable thin provisioning on a container:

1. Select the NAS container.
2. Click **Modify Settings**.
3. Click the **Space** tab.
4. Select the **Thin Provision** checkbox.
5. Move the slider (arrow below gray bar) to adjust the space reservation values.
6. Click **OK**.

NOTE:

- **The NAS clusterwide default values are applied when new containers are created.**
- **If you select the Container thin provisioning default checkbox, ensure that you click the Save all changes icon afterward. If you do not, updates from the array will clear your selection.**

Thin-Provisioned Container Attributes

Thin-provisioned containers have the following characteristics:

- Minimum reserve is 0% to 99%. The default is 0%.
- Maximum size for thin-provisioned container is 500TB.
- The amount of space reserved for a container is either the minimum reserve or the actual used space, whichever is larger.
- Space consumption grows linearly with used space when usage exceeds the minimum reserve.

Disable Thin Provisioning on a Container

To disable thin provisioning on a container:

1. Select the NAS container.
2. Click **Modify Settings**.
3. Click the **Space** tab.
4. Clear the **Thin Provision** checkbox. The space reservation slider is now cleared.
5. Click **OK**.

About Data Rehydration

Rehydration is the process whereby data that has been compressed becomes decompressed and readily accessible. Rehydration happens automatically and occurs as necessary to provide access to compressed data. Whenever data is vacated from one group member to another, any compressed data is automatically rehydrated during that process.

NOTE: Rehydrating data affects performance and disk storage space utilization, because the data blocks are rendered on the disk in their original form.

When you disable data reduction, data remains in its reduced state during subsequent read operations by default. You can also enable rehydrate-on-read when you disable data reduction. Enabling rehydrate-on-read causes a rehydration (a reversal of data reduction) of data on subsequent read operations. You cannot rehydrate an entire NAS container unless you read the entire container.

NOTE: Neither data reduction nor data rehydration occurs on containers with less than 5GB of free space.

Rehydrate on Read with User Quotas

If you defined user quotas based on logical space utilization, those quotas remain in force throughout rehydration of the reduced data. For example, if three users have stored identical files of 3GB each, then after applying data reduction, only one physical copy will remain, thus consuming 3GB rather than 9GB. If each user has a user quota set to 4GB, then each of them can store only 1GB of additional data.

Enable Rehydrate on Read

NOTE: You must be logged in as grpadmin to enable rehydrate-on-read on an existing NAS container.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Modify settings** to open the dialog box.
3. Click the **Data Reduction** tab.
4. Clear the **Enable data reduction** checkbox to display the Rehydrate on read option.
5. Select the **Rehydrate on read** checkbox, and click **OK**.
6. In the confirmation window, click **Yes**.

The data reduction settings are shown at the bottom of the status information panel.

NAS Container Data Reduction

Data reduction is a process that runs according to a schedule on each NAS container that has data reduction enabled. A policy that you define determines whether or not a file qualifies for data reduction, on the basis of access and modification times of that file. The data reduction process analyzes files to determine if multiple copies of an individual file, or portions of a file, or blocks, can be saved in a more efficient format.

The PS Series controller manages the data reduction schedule, which defines when the data reduction process begins and ends, and how frequently data reduction is performed. The NAS array controller manages the data reduction process.

NOTE: To configure and enable the data reduction feature on a NAS container, you must have group administrator (grpadmin) privileges.

User quotas are enforced based on logical data size, not on the physical footprint of the data. If two users each have a copy of the same 1GB file, the physical footprint of the file is 1GB after data reduction, and both users use 1GB of their individual user quota. If data reduction is disabled, the user quota will continue to be enforced based on the logical space utilization until all the reduced data is rehydrated.

Data reduction occurs:

- During the time and days specified in a schedule.
- When a file has not been accessed or modified for the duration of time defined by the file filters in the data reduction policy.

NOTE: Enabling data reduction on a container permanently removes the snapshot reserve functionality for that container. Snapshot reserves cannot be enabled even if you later disable data reduction. Snapshots can still be manually created, deleted, and restored, and automatically deleted when the maximum number of snapshots to keep is exceeded.

Data reduction is accomplished through two methods:

- Deduplication — Applied whenever data reduction is enabled. Deduplication replaces qualified duplicated data with a pointer to a single copy of the data.
- Compression — Applied only when compression is enabled as part of the data reduction policy. Compression occurs within files when individual blocks can be saved in a more efficient format than originally written.

Data reduction is performed on blocks in 128KB chunks. Space savings from data reduction are reported in MB units.

Deleting a reduced file will not result in additional free disk space on the container, unless the file being deleted is the last copy of that file. Free space gained from deleting the last copy of a file will not be available until the disk cleaner process cleans the file marked for deletion.

When you enable data reduction on a NAS container, default data reduction policy settings are applied to that container. You can modify the default policy settings or define other policy settings. The data reduction policy specifies if data compression is applied and determines if files qualify for data reduction based on the last time they were accessed and modified. For more information about setting default values, see Creating Default Data Reduction Properties of a NAS Cluster.

Data reduction affects system performance. When a reduced file is accessed, the controller reads a copy of the deduplicated data and uncompresses the data as it is being sent to a client. Unless the reduced data is modified, it remains reduced on the disk. After data reduction is disabled, if at least 5GB of free space is available on the container, when data is read on the container, the data is uncompressed.

When deciding whether or not to enable data reduction, consider the storage resource and data reduction processing efficiency of your system.

Data Reduction Methods

Data reduction is supported on a per-NAS-volume basis to store data more efficiently. The Dell FluidFS cluster supports two types of data reduction:

- Deduplication — Performed on qualified data when data reduction is enabled on a container. You cannot disable deduplication when data reduction is enabled. Deduplication (or *dedupe*) provides data reduction by eliminating redundant copies of data across files in a volume by keeping only one copy of unique deduplicated data. When the data is read, it is uncompressed and rehydrated in memory to its original form. The data on the disk remains unchanged.

  NOTE: Deduplication is enabled by default. It activates *when* you enable data reduction, because it is part of the data reduction process. The `Deduplication` option shows as being enabled even when `Enable data reduction` is not yet selected.
- Compression — Uses a Level Zero Processing System (LZPS) compression algorithm on already deduplicated data to further enhance data reduction. The next time the data is accessed, it is uncompressed and rehydrated, thereby returning to its original form. (To utilize this feature, when you create the NAS volume, first enable data reduction and then enable compression.)

Data Reduction and Snapshot Reserve

The snapshot reserve functionality enables you to reserve a maximum percentage of disk storage space for snapshots. When a snapshot is created, it serves as a pointer to the data contained within the NAS container. It does not write to or consume any space in the snapshot reserve until the data is modified or deleted from the container.

The oldest snapshots autodelete in the following instances:

- The data written to the snapshot reserve exceeds the amount of reserved space.
- The number of scheduled snapshots exceeds the specified maximum number of snapshots to keep (max-keep).

The max-keep guideline does not apply to manually created snapshots, only to scheduled snapshots. Manually created snapshots can coexist with scheduled snapshots. They can also be autodeleted when the snapshot reserve space consumed exceeds the amount of reserved space if data reduction is not enabled.

The snapshot reserve functionality is permanently disabled after data reduction is enabled, even if data reduction is later disabled on the container. When you enable data reduction on a container, the only way to automatically delete snapshots when running a snapshot schedule is to define a maximum number of snapshots to retain for that schedule.

All other snapshot actions are supported after data reduction is enabled, as shown in Table 47. Data Reduction Effect on Snapshots.

Table 47. Data Reduction Effect on Snapshots

Supported Snapshot Actions	Before Data Reduction	After Data Reduction
Manually create snapshots	Yes	Yes
Manually delete snapshots	Yes	Yes
Automatically create snapshots using snapshot schedules	Yes	Yes
Automatically delete snapshots when a maximum number of snapshots is exceeded	Yes	Yes
Automatically delete snapshots when a maximum percentage of container space (snapshot reserve) is exceeded	Yes	No
Yes: Enabled No: Disabled		

Enable Data Reduction

You must be logged in as grpadmin to enable data reduction on an existing NAS container.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Modify settings** to open the dialog box.

3. Click the **Data Reduction** tab.
4. Select the **Enable data reduction** checkbox. A confirmation message is displayed.
5. Click **Yes** to confirm enabling data reduction.
6. Review the data reduction policy shown on the dialog box.
7. (Optional) If you want to modify the policy settings for the container, click the **Modify policy** button to open the dialog box.

   - Compression

     By default, compression is disabled and only deduplication is enabled.
   - Access Time

     The access time is the minimum number of days that must pass since the file was last accessed before the file is eligible for data reduction. This value must be in the range of 30 to 365 if compression is enabled. If compression is disabled, the range is 5 to 365 days, inclusive. By default, the value for both times is 30 days.
   - Modify Time

     The modify time is the minimum number of days that must pass since the file was last modified before the file is eligible for data reduction. This value must be in the range of 30 to 365 if compression is enabled. If compression is disabled, the range is 5 to 365 days, inclusive. By default, the value for both times is 30 days. By default, the value for both times is 30 days.
8. Click **OK**.
9. In the Modify Settings dialog box, click **OK**. A confirmation message is displayed.
10. Click **Yes** to apply the container settings.

NOTE: The status of the filter is listed as `File filters`….`.disabled` if the ignore-filters option was set through the CLI. If the status is `File filters`…`disabled`, any filters that have been configured through the GUI (or CLI) for Access Time or Modify Time have been disabled and all files are candidates for data reduction. Specifying the `ignore-filters` flag enables data reduction on a container with archive data without waiting for the minimum Access Time/ Modify Time data-reduction policy. Dell recommends that you not use this setting on containers with live data. That is, you should use this setting only on containers with archive data.

Modify NAS Container Data Reduction Settings

NOTE: To change data reduction settings of an existing NAS container, you must have group administrator (grpadmin) privileges.

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Modify settings**.
3. In the **Modify settings** dialog box, type or select the appropriate data.

   For information about the fields and description in each tab, see Create a NAS Container.

Modify NAS Cluster Default Data Reduction Settings

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. Click the **Modify** link to change data reduction schedule settings.
4. If you want to perform compression as part of data reduction, select **Enable compression**.
5. (Optional) To change the access time policy, enter a new value in **Access Time**. After modifying the access time, click the **Save** icon to save the changes. The access time is the minimum number of days that must pass since the file was last accessed before the file is eligible for data reduction. This value must be in the range of 30 to 365 if compression is enabled. If compression is disabled, the range is 5 to 365 days, inclusive. By default, this value is 30 days.
6. (Optional) To change the modification policy, enter a new value in **Modify Time**. After modifying the modify time, click the **Save** icon to save the changes. The modify time is the minimum number of days that must pass since the file was last modified before the file is eligible for data reduction. This value must be in the range of 30 to 365 if compression is enabled. If compression is disabled, the range is 5 to 365 days, inclusive. By default, this value is 30 days.
7. Click the Save all changes icon in the toolbar.

NOTE: The status of the filter is listed as `File filters….disabled` if the ignore-filters option was set through the CLI. If the status is `File filters…disabled`, any filters that have been configured through the GUI (or CLI) for Access Time or Modify Time have been disabled and all files are candidates for data reduction. Specifying the `ignore-filters` flag enables data reduction on a container with archive data without waiting for the minimum Access Time/ Modify Time data-reduction policy. Dell recommends that you not use this setting on containers with live data. That is, you should use this setting only on containers with archive data.

Data Reduction Policy

A data reduction policy defines which files qualify for data reduction and what type of data reduction should be performed.

The policy is based on:

- Access Time — The last day that the file was open for reading. (The filter is based on the number of days and not a specific time.)
- Modification Time — The last day that the file was changed. (The filter is based on the number of days and not a specific time.)
- Compression — Perform file compression, in addition to deduplication. Use compression with text files to increase storage space efficiency. Compression does not increase storage space efficiency with binary files, such as audio and video files. Deduplication, which is enabled whenever data reduction is enabled, increases disk storage space utilization efficiency with both text and binary files.

Modify Container Data Reduction Policy

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Modify settings** to open the dialog box.
3. Click the **Data Reduction** tab.
4. Select the **Enable data reduction** checkbox.
5. Click the **Modify policy** button.
6. (Optional) To perform compression as well as deduplication, select the **Enable compression** checkbox.
7. To change the access time required for data reduction, enter a new value in **Access Time**. After modifying the access time, click the Save icon to save the changes. The access time is the minimum number of days that must pass since the file was last accessed before the file is eligible for data reduction. This value must be in the range of 30 to 365 days if compression is enabled. If compression is disabled, the range is 5 to 365 days. By default, this value is 30 days.
8. To change the modification time required for data reduction, enter a new value in **Modify Time**. After modifying the modify time, click the Save icon to save the changes. The modify time is the minimum number of days that must pass since the file was last modified before the file is eligible for data reduction. This value must be in the range of 30 to 365 days if compression is enabled. If compression is disabled, the range is 5 to 365 days. By default, this value is 30 days.
9. Click **OK**.

NOTE: The status of the filter is listed as `File filters….disabled` if the `ignore-filters` option was set through the CLI. If the status is `File filters…disabled`, any filters that have been configured through the GUI (or CLI) for Access Time or Modify Time have been disabled and all files are candidates for data reduction. Specifying the `ignore-filters` flag enables data reduction on a container with archive data without waiting for the minimum Access Time/ Modify Time data-reduction policy. Dell recommends that you not use this setting on containers with live data. That is, you should use this setting only on containers with archive data.

Modify Default Data Reduction Policy

The default data reduction policy sets the access time and modify time values that are used by default when you enable data reduction on a container.

To modify the default data reduction policy:

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. Modify the settings on the Default Data Reduction Policy panel as needed:
   - In **Access Time**, enter a value. After modifying the access time, click the Save icon to save the changes. The access time is the minimum number of days that must pass since the file was last accessed before the file is eligible for data reduction.

This value must be in the range of 30 to 365 days if compression is enabled. If compression is disabled, the range is 5 to 365 days. By default, this value is 30 days.

- In **Modify Time**, enter a value. After modifying the modify time, click the Save icon to save the changes. The modify time is the minimum number of days that must pass since the file was last modified before the file is eligible for data reduction. This value must be in the range of 30 to 365 days if compression is enabled. If compression is disabled, the range is 5 to 365 days. By default, this value is 30 days.

4. Click **OK** to save your changes.
5. Click the Save all changes icon to save your changes.

NOTE: The status of the filter is listed as `File filters….disabled` if the ignore-filters option was set through the CLI. If the status is `File filters…disabled`, any filters that have been configured through the GUI (or CLI) for Access Time or Modify Time have been disabled and all files are candidates for data reduction. Specifying the `ignore-filters` flag enables data reduction on a container with archive data without waiting for the minimum Access Time/ Modify Time data-reduction policy. Dell recommends that you not use this setting on containers with live data. That is, you should use this setting only on containers with archive data.

Create Default Data Reduction Properties

NOTE: To enable or disable data reduction properties at a cluster level, you must have group administrator (grpadmin) privileges.

To set default data reduction properties at a NAS cluster level:

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. In the **Default Data Reduction Policy** panel, type or select the necessary data.

About NAS Data Reduction Schedules

Data reduction occurs as a scheduled activity. A data reduction schedule is defined at the NAS cluster level and controls when data reduction is run on all the containers that have data reduction enabled.

You can define multiple data reduction schedules at the NAS cluster level. If two or more data reduction schedules overlap, data reduction begins when the first schedule begins running and ends when the last schedule completes running. You can define up to a total of 1024 schedules on each NAS cluster. This total includes all defined data reduction, NAS, snapshot, and replication schedules.

You can configure a data reduction schedule to define the days and times that data reduction begins and ends. If you do not change the default data reduction schedule settings, data reduction will run every day, starting at midnight and running until 6:00 a.m.

When you define a data reduction schedule:

- Specify start times and end times in a 12-hour format (such as, 3:30PM, 7AM).
- Schedule data reduction during off-peak hours to avoid potential performance impact.

When multiple containers have data reduction enabled, data reduction processes qualified files for a fixed amount of time starting with the first container. Processing continues on qualified files within the next container, and so on, until data reduction processes all qualified files on each of the enabled containers. If data reduction does not complete on all the enabled containers when the schedule ends, data reduction continues to process the remaining data when the schedule next runs.

NOTE: To avoid any potential performance impact from running a data reduction process, define a schedule to run data reduction during off-peak hours.

For example:

Data reduction is enabled to run on the two containers, NAS_A and NAS_B, of a NAS cluster that is managed by one controller. When the schedule finishes running, if data reduction is being performed on container NAS_B, when the data reduction schedule begins running the next time, data reduction begins running on container NAS_B. When data reduction completes on NAS_B, the data reduction service begins running on container NAS_A.

Analyzing data on a container for data reduction will take longer than scanning and reducing previously analyzed data. When defining the data reduction schedule on your system, consider when all files on a container need to have been analyzed for data reduction. The efficiency of the data reduction process is affected by the number of controllers analyzing data.

Create a Data Reduction Schedule

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. In Data Reduction Schedules, click the **Add** link. The Create Data Reduction Schedule dialog box opens.
4. In **Name**, type a name for the new schedule. The name can be up to 63 characters long.
5. In **Days**, select the days that the schedule will run. You must select at least one day.
6. In **Start Time**, specify in 12-hour format the time the schedule will start. For example, 3:30PM or 7:00AM.
7. In **End Time**, specify in 12-hour format the time the schedule will stop.
8. If you are ready to run data reduction on the container, select the **Enable schedule** checkbox.

   If you do not enable the schedule now, you can enable the schedule later from the Data Reduction Schedules panel.
9. Click **OK**.

Update a NAS Cluster Data Reduction Schedule

NOTE: To edit the properties of a NAS cluster data reduction schedule, you must have group administrator (grpadmin) privileges.

1. Cick **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. In the Data Reduction Schedules panel, select the schedule and then click **Modify**.
4. In the **Modify data reduction schedule** dialog box, type or select the necessary data.
5. Click **OK**.

   The updated information is saved and becomes effective immediately.

Modify a Data Reduction Schedule

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. On the Data Reduction Schedules panel, select the schedule that you want to modify.
4. Click **Modify**. The Modify Schedule dialog box opens.
5. Modify the schedule as appropriate:

   a. In **Name**, type a descriptive name for the schedule.
   b. In **Days**, select each day that you want to schedule to run. You must select at least one day.
   c. In **Start Time**, specify the time that the schedule will start running.
   d. In **End Time**, specify the time that the schedule will stop running.
   e. In Schedule Status, select or clear the **Enable schedule** checkbox. A schedule that is not enabled will not run.
6. Click **OK**.

Delete a Data Reduction Schedule

1. Click **Group**, expand **Group Configuration**, and select a NAS cluster.
2. Click the **Data Reduction** tab.
3. In the Data Reduction panel, select the schedule that you want to delete and click **Delete**. The Delete Schedule dialog box opens.
4. Click **Yes**.

FS Series VAAI Plugin

The VAAI plugin allows ESXi hosts to offload some specific storage-related tasks to the underlying FluidFS appliances. The plugin supports the following VAAI NAS Primitives:

- **Full File Clone**– Offload the creation of a virtual disk full clone
- **Fast File Clone** (Native Snapshot) – Offload the creation of a virtual disk linked clone
- **Extended Statistics** – Query for space usage on FS series datastores

Installing the plugin enables VAAI NAS primitives for all datastores residing on FS Series v4 or later systems, adding the following functionalities:

1. Virtual machine cloning from vCenter will request FS Series appliances to generate a full copy of the corresponding machine.
2. The creation of virtual machine linked clones will be offloaded to FS series appliances.

The plugin is provided in two alternate forms. Both forms can be downloaded from the FTP server ftp://*<FluidFS_Cluster_public IP>*:44421/vaai_plugin:

- A VIB file – **FluidFSNASVAAI_For_Esx_v5.5.vib** file
- A depot – **FluidFSNASVAAI_For_Esx_v5.5.zip** file

Both forms provide equal functionality, differing only in the way they are installed.

Installation Instructions

The FS Series VAAI plugin supports ESXi versions 5.5, 5.5U1, and 5.5U2.

NOTE: The FS Series VAAI plugin should be installed on each relevant ESXi host and requires a reboot.

1. Connect to FS Series via FTP on port 44421 using administrative credentials.
2. Download the VAAI plugin file located inside the /vaai_plugin folder.
3. Transfer the file to the /tmp/ folder of the ESXi host.
4. Install the plugin, depending on the file type that you transferred:

   -
     ```
     ~ # esxcli software vib install -d /tmp/FluidFSNASVAAI_For_Esx_v5.5.zip
     ```

   or

   -
     ```
     ~ # esxcli software vib install –v esxcli software vib install -v file:///tmp/
     FluidFSNASVAAI_For_Esx_v5.5.vib
     ```
5. Reboot the ESXi host.

Plugin Verification

To check if the VAAI plugin is installed in an ESXi host, type the following command in the ESXi console:`# esxcli software vib list | grep Dell_FluidFSNASVAAI`

A positive reply should return:

```
Dell_FluidFSNASVAAI 1.1.0-250 DELL VMwareAccepted 2015-02-17
```

To verify that an FS Series datastore has VAAI enabled use the command vmkfstools –P in the ESXi host console. The following example illustrates the query and output for a datastore named FSseries_datastore residing on a FS Series v4 or later system:

```
~ # vmkfstools -Ph /vmfs/volumes/FSseries_Datastore/

NFS-1.00 file system spanning 1 partitions

File system label (if any): FSseries_Datastore

Mode: public

Capacity 200 GB, 178.3 GB available, file block size 4 KB, max file size 16777216 TB

UUID: 1cec81cb-6db87d1c-0000-000000000000

Partitions spanned (on "notDCS"):

        nfs:FSseries_Datastore

NAS VAAI Supported: YES

Is Native Snapshot Capable: YES
```

Removal Instructions

To remove the VAAI plugin from an ESXi host:

1. Execute the following command in the ESXi host console:
   ```
   ~ # esxcli software vib remove -n Dell_FluidFSNASVAAI
   ```
2. Reboot the ESXi host.

15

Diagnose and Resolve NAS Cluster and PS Series Issues

If you have to work with Dell support to resolve an issue related to a PS Series array or a NAS cluster, you can provide the support team with necessary data to facilitate successful troubleshooting of the issue without having to install software or download tools from the Dell support site. You can send the report from an alternate system and be certain that the data collected from your server is not viewable by nonauthorized individuals during the transmission to Dell support.

NAS Cluster Diagnostics

To diagnose and solve a problem on a NAS cluster, you must first determine if the problem is hardware- or software-related. It is important to replace failed hardware as soon as possible. If a controller that is part of an intact NAS controller pair becomes unavailable (for example, if the controller has no power, has failed, or is disconnected from the network), the controller pair is degraded.

Although the NAS cluster is still operational, you cannot perform most service modifications until you detach the failed controller from the controller pair. It is important to determine if the problem resides in the PS series group, the NAS controllers, the network configuration, or the client configuration. The following tasks are useful:

- Review the PS Series group event log, which contains events related to a NAS cluster, including the service itself, the NAS controllers, NAS containers, SMB shares, and NFS exports.
- Be aware of the NAS cluster alarms, which help you to identify and correct problems before any operations are affected.
- Contact Dell support for assistance in running and interpreting online or offline diagnostics, which can provide valuable information.

About Current and Past NAS Network Performance Statistics

You can use network performance statistics to help with:

- Evaluating bottlenecks
- Troubleshooting
- Scheduling maintenance
- Monitoring performance

You can gather current and past network performance statistics by using the CLI.

Statistics are available in the following categories:

1. Cluster Network
2. Controller Network
3. Controller Load-Balancing (provides data that is a summary of all available statistics)

For each category of statistics, you can choose from the following time periods:

1. `Current` – Data is shown for a one-time interval.

2. `Day` – Last 24 hours. Data is shown for each hour.
3. `Week` – Last 7 days. Data is shown every 6 hours in the last 7 days.
4. `Month` – Last month. Daily for the last month.
5. `Year-` Last year. Data is shown for every 2 weeks of the last year.

Online Diagnostics

You can obtain online diagnostic information while the system is still online and running.

The following diagnostic options are available from the Group Manager CLI:

- Performance
- Network
- Protocols
- NAS containers
- General System

For more information, see the *Dell EqualLogic Group Manager CLI Reference Guide*.

Offline Diagnostics

To obtain offline diagnostics, the NAS cluster must be offline or not serving data. Diagnostics can help you troubleshoot low-level hardware issues.

Offline diagnostics use the following Dell tools:

- Dell Diagnostics Utility

  This tool helps you check your computer's hardware without any additional equipment or destroying any data. If you find a problem that you cannot solve by yourself, the diagnostic tests can provide important information that you will need when talking to Dell service and support personnel.

  NOTE: Use the Dell Diagnostics Utility only to test Dell systems. Using this program with other systems might cause incorrect system responses or error messages.

- MP Memory

  A test tool developed by Dell and based on MS-DOS, it is efficient for configurations that have more than 4GB of memory. The tool supports single-processor or multiprocessor configurations, as well as processors using the Intel hyper-threading technology. MP Memory operates only on Dell PowerVault servers that are based on Intel processors.

  This tool complements Dell 32-bit diagnostic utility tests and helps provide complete, comprehensive diagnostics on the NAS controller.

Run Offline Diagnostics

To run offline diagnostics:

1. Insert the Dell EQL FS7500 Resource media into the controller's DVD drive and reboot the controller. The NAS controller boots to the DVD.
2. From the FS7500 Restore and Diagnostic Utilities menu, select **Hardware Diagnostics**.

   CAUTION: Do not select either the File System Reinstall or Firmware Reset option. The first option reinstalls the image on your NAS controller and might cause the loss of data. The second option resets IP addresses on the NAS controller and also might cause the loss of data.
3. The **Hardware Diagnostics** menu displays the following choices:

```
1 Mpmemory diagnostic (supports console-redirection in output log only).
2 Delldiag text-based diagnostic (Full console-redirection support).

3 Loop Mpmemory and diagnostic in batch mode.
  *** Please install all removable media if selecting option 3.
4 Quit
Enter option or letter: (default = 3, timeout in 16 secs)
```

4. Choose the appropriate option.
5. Press the Escape key at any time during a test to stop the test.

Generating PS Series and NAS Cluster Diagnostics Reports

NOTE: To generate diagnostics reports, you must have group administrator (grpadmin) privileges.

To generate a diagnostics report:

1. Log in to Group Manager by using your administrator login ID and password.
2. In the navigation pane, click **Tools**.
3. Click **Diagnostics reports**.
4. In the **Generate and email diagnostics** wizard, type or select the necessary data in the pages.

Adding SMTP Servers

1. Log in to Group Manager by using your administrator login ID and password.
2. In the navigation pane, click the **Tools** tab.
3. Click **Diagnostics reports**.
4. In the **Generate and email diagnostics** wizard, type or select the necessary data in the pages.
5. In the PS Series Diagnostics–Report Destination wizard, click **Configure SMTP**.
6. In the Configure Email dialog box, click **Add**.
7. In the **Add SMTP server** dialog box, type the IP address of the SMTP server.
8. For **Port**, type the port number at which the network connection must be started.

   To use the default port, click **use default port**.
9. Click **OK**.

   The SMTP server is added and listed in **SMTP servers**.

Modifying SMTP Server Properties

1. Log in to Group Manager by using your administrator login ID and password.
2. In the navigation pane, click **Tools**.
3. Click **Diagnostics reports**.
4. In the **Generate and email diagnostics** wizard, type or select the necessary data in the pages.
5. In the PS Series Diagnostics–Report Destination wizard, click **Configure SMTP**.
6. In the Configure Email dialog box, under the SMTP servers list, select the SMTP server and then click **Modify**.
7. In the **Modify SMTP server** dialog box, edit the IP address of the SMTP server.
8. In **Port**, edit the port number to which the connected must be started.
9. Click **OK**.

   The SMTP server properties are updated and listed in **SMTP servers**.

Deleting SMTP Servers

1. Log in to Group Manager by using your administrator login ID and password.
2. In the navigation pane, click **Tools**.
3. Click **Diagnostics reports**.
4. In the **Generate and email diagnostics** wizard, type or select the necessary data in the pages.
5. In the PS Series Diagnostics–Report Destination wizard, click **Configure SMTP**.
6. In the Configure Email dialog box, select the SMTP server and then click **Delete**.

The following message is displayed:

`Do you want to delete` *`TCP/SCP_server_ip_address`*`?`

7. Click **Yes**.

   The server is deleted from the list and the SMTP servers section is updated.

Reinstall FS Series Firmware v4 from an Internal USB

FS7600 and FS7610 NAS appliances contain an internal USB from which you can reinstall the FS Series firmware v4 factory image. If you experience general system instability or a failure to boot, you might have to reinstall the image.

CAUTION: Reinstalling the NAS cluster software reverts your system to factory defaults. All data on the NAS system will be unrecoverable after performing this procedure.

NOTE: After reinstalling the NAS system software, also install the latest service pack updates.

Before reinstalling FS Series firmware v4:

- If your controller is still an active member in the NAS cluster, you must first detach it using the Group Manager GUI or CLI.
- Connect a keyboard and monitor to the appliance.
- Back up your data (recommended by Dell).

To reinstall FS Series firmware v4:

1. Power off the controller using the recessed power button on the back panel of the FS7600 or FS7610 NAS appliance. For information about powering on the appliance, see the *Hardware Owner's Manual* for your system.

   NOTE: Power off only the controller on which you are reinstalling FS Series firmware v4. Do not power off the second controller, because the second controller functions normally in journaling mode while the other controller is down.
2. After the controller powers off, power it back on using the recessed power button.
3. After the BIOS starts, when prompted to enter the boot menu, press F11. You might have to press F11 a few times to access the menu
4. Select **Generic STORAGE DEVICE**.
5. On the next screen, select **FluidFS Reinstall**.
6. Select the version of FS Series firmware to install, either v4.0.x or 3.0.940200. Use option 1 with PS Series firmware version 8.1 or later and option 2 with PS Series firmware version 8.0 or later. See Figure 14. Dell FluidFS Reinstallation Menu

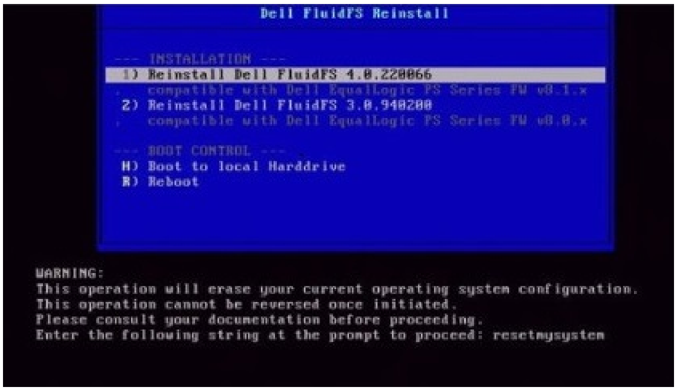

```
                    Dell FluidFS Reinstall

    --- INSTALLATION ---
     1) Reinstall Dell FluidFS 4.0.220066
    .   compatible with Dell EqualLogic PS Series FW v8.1.x
     2) Reinstall Dell FluidFS 3.0.940200
    .   compatible with Dell EqualLogic PS Series FW v8.0.x

    --- BOOT CONTROL ---
     H) Boot to local Harddrive
     R) Reboot

WARNING:
This operation will erase your current operating system configuration.
This operation cannot be reversed once initiated.
Please consult your documentation before proceeding.
Enter the following string at the prompt to proceed: resetmysystem
```

Figure 14. Dell FluidFS Reinstallation Menu

7. When prompted, confirm the action by typing `resetmysystem` and pressing Enter. .

After the reinstallation completes, the controller reboots into standby mode.

After reinstalling FS Series firmware v4.0, use the Group Manager GUI or CLI to attach the controller to a NAS cluster.

16

About Backing Up and Protecting Your Data

A PS Series group is part of a comprehensive backup and data protection solution.

Snapshots provide quick recovery and offloading backup operations. On a PS Series group, the system creates the copy instantly and maintains it on disk storage within the group. It does not disrupt access to the volume and requires minimal impact on running applications. Snapshots can provide a stable copy of data for copying to backup media. Restore operations are more reliable because snapshots ensure the integrity of the backed-up data.

Cloning a volume creates a duplicate volume with a new name and iSCSI target. You can then create a writable volume.

Replication protects data from serious failures such as destruction of a volume during a power outage, or a complete site disaster. Asynchronous replication remotely replicates data from one PS Series group to another over a standard IP network over long distances. You can quickly configure volumes for replication. The PS Series arrays manage the underlying hardware resource complexity. Synchronous replication allows replication across two different storage pools in the same PS Series group, resulting in two hardware independent copies of the volume data.

You can create schedules to automatically perform volume and NAS container operations at a specific time or on a regular basis (for example, hourly or daily).

The Manual Transfer Utility (MTU) allows you to export a volume on to a transportable media so you can send or carry the media to a remote site.

About Volume Data Protection

Dell recommends that you use snapshot and replication functionality to protect volume data.

A snapshot is a point-in-time representation of a PS Series iSCSI volume. Seen on the network as an iSCSI target, this snapshot is maintained in an array as deltas from the original volume. Snapshots can protect against human error, viruses, or database corruption. You can recover data from a snapshot by setting it online or by restoring the volume from a snapshot.

To protect against disasters, you can replicate volume data from one group to another. A replica is a point-in-time representation of a PS Series iSCSI volume. The original volume and its replica are located on different PS Series groups (replication partners) potentially separated by some geographical distance to facilitate disaster tolerance. The replicated volume that is located on a different group does not depend on the original volume so that, in the event of a disaster, you can host the volume from the recovery group and later fail back to the original group with minimal disruption to users.

Synchronous replication (SyncRep) is the simultaneous writing of volume data across two different storage pools in the same PS Series group, resulting in two hardware-independent copies of volume data. Each write must go to both pools before the write is acknowledged as complete. If one pool is not available, you can obtain volume data from the other pool.

Dell recommends that you protect data by using a robust backup application as a precaution in the event of multiple drive failures, in addition to snapshot and replication functionality.

Protect NAS Container Data with NDMP

A NAS cluster supports the Network Data Management Protocol (NDMP), which facilitates backup operations for network-attached storage, including NAS containers.

A NAS cluster includes an NDMP server that performs NAS container backups to an external Data Management Application (DMA) server running backup software.

After you configure a DMA server for a NAS cluster, the NDMP server listens on the client network for backup requests from the DMA servers. You can configure up to 10 DMA servers.

The DMA server then accesses, or mounts, the NAS containers that it wants to back up and initiates the backup operations. For example, to back up NAS container data:

1. The DMA server creates a connection to the NAS cluster IP address.
2. The NDMP server on the NAS cluster creates a temporary snapshot of each NAS container that the DMA server designated for backup. The snapshot reserve for the NAS container stores the temporary snapshot.
3. The NDMP server copies the NAS container data to the DMA server.
4. After receiving the data, the DMA server moves the data to a storage device, such as a local disk or tape device.
5. When the backup completes, the NDMP server deletes the temporary snapshots.

NOTE: Manually deleting the temporary NDMP snapshot immediately terminates the current backup session. If you change the NAS cluster IP address, the NDMP service restarts. In addition, if you change the NAS cluster IP address, you must also make the reciprocal change on the DMA servers.

Configure NDMP for a NAS Cluster

You must have snapshot reserve to use the Network Data Management Protocol (NDMP) to facilitate NAS container backup operations.

NOTE: Select the Enable data reduction checkbox to activate the Modify policy button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.

Configure NDMP Settings

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Advanced** tab.
3. In the NDMP panel, click **Modify Settings**.
4. In the **Modify NDMP settings** dialog box, type information in the various fields (user name, password, confirm password, port).

   NOTE: NDMP passwords accept up to 31 ASCII characters, spaces, and punctuation (except quotation marks).

5. Click **OK**.

Configure DMA Servers

1. Click **Group**, expand **Group Configuration**, and then select the NAS cluster.
2. Click the **Advanced** tab.
3. In the NDMP panel, click **Configure DMA Servers**.
4. In the Configure DMA Servers dialog box, click **Add**.
5. Type the IP address for a Data Management Application (DMA) server. You can add up to 10 DMA servers.
6. Click **OK**. The DMA server appears in the list.
7. Click **OK** to dismiss the Configure DMA Servers dialog box.

About Snapshots

Snapshots enable you to capture volume data at a specific point in time without disrupting access to the volume.

A snapshot represents the contents of a volume at the time of creation. You can create snapshots of standard volumes, in addition to template volumes and thin clone volumes. A volume can be restored from a snapshot or a snapshot can be cloned to create a new volume.

Creating a snapshot does not prevent access to a volume, and the snapshot is instantly available to authorized iSCSI initiators. Similar to volumes, snapshots appear on the network as iSCSI targets, and can be set online and accessed by hosts with iSCSI initiators.

You can create a snapshot of a volume at the current time, or you can set up schedules to automatically create snapshots on a regular basis.

If you accidentally delete data, you can set a snapshot online and retrieve the data. If a volume is corrupted, you can restore the volume from a snapshot. You can also clone a snapshot to create a new copy of a volume.

NOTE:

- **You cannot restore a template volume from a snapshot.**
- **Generally, snapshots will not be deleted unless you take action to delete them. In some instances, however, snapshots can be deleted by the system. For example, when a new snapshot is taken and not enough snapshot reserve space is available for the new snapshot and the previous one, the older one will be deleted. A snapshot can also be deleted during snapshot borrowing if you run out of borrowable space.**
- **Ensure that no NAS snapshots have these internal reserved names:**
  - **rep_***
  - **rollback**
  - **current**

How Snapshots Work

Snapshots simplify and increase the performance of backup and recovery operations. Snapshots enable you to capture volume data at a specific point in time without disrupting access to the volume.

To create snapshots of a volume, you must allocate snapshot reserve for the volume. Initially, a snapshot consumes no space from the snapshot reserve because it shares all data with the volume (sometimes called the base volume). When the volume changes, the snapshot reserve tracks those changes to maintain the volume contents at the time of snapshot creation.

NOTE:

- **When you delete a volume, the group deletes all snapshots of that volume.**
- **Select the Enable data reduction checkbox to activate the Modify policy button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.**

As with volumes, snapshots appear on the network as iSCSI targets. All the iSCSI target security mechanisms apply to snapshots.

You can access the data in a snapshot by using the following methods:

- Restore a volume from a snapshot – This operation replaces the volume with the data that existed at the time you created the snapshot.
- Clone a snapshot – The new volume contains the volume data that existed at the time you created the snapshot.
- Set the snapshot online – iSCSI initiators can access the target in the usual way.

You can create snapshots of individual volumes or volume collections. When you perform a snapshot operation on a volume collection, the group creates a set of snapshots (one for each volume in the collection) called a snapshot collection. You can also simultaneously create snapshots of multiple volumes that are not in a collection. The resulting set of snapshots is called a custom snapshot collection.

Use volume schedules to create snapshots at a specific time or at a time interval (such as hourly, daily, or weekly).

NOTE: Generally, snapshots will not be deleted unless you take action to delete them. In some instances, however, snapshots can be deleted by the system. For example, when a new snapshot is taken and not enough snapshot reserve space is available for the new snapshot and the previous one, the older one will be deleted.

About Snapshot Reserve

Before you can create snapshots of a volume, you must allocate snapshot reserve for the volume. Snapshot reserve is consumed from the pool where the volume resides.

You can allocate snapshot reserve when you create a volume, or you can modify a volume's properties to change the snapshot reserve.

Snapshot reserve is a percentage of the volume reserve. Because the volume reserve for a thin-provisioned volume changes as volume usage increases, the snapshot reserve for a thin-provisioned volume also changes.

The group generates event messages when the amount of free snapshot reserve falls below a user-defined threshold.

NOTE:

- **Depending on the policy that you set for snapshot space recovery, the group preserves snapshot reserve.**
- **Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.**

Snapshot Reserve Settings

The following snapshot reserve settings use groupwide default values, unless you explicitly change them for a volume:

- Snapshot reserve

  Amount of space, based on a percentage of the volume reserve, that the group allocates to snapshots. When you create a volume, you can specify the snapshot reserve percentage for the volume. Otherwise, the group applies the groupwide default value. You can modify the snapshot reserve value as needed in **Group Manager → Group Configuration** , under the Default tab.
- Snapshot space recovery policy

  Action the group takes when a new snapshot exceeds snapshot reserve:

  – Delete the oldest snapshots to free space for new snapshots
  – Set the volume (and snapshots) offline
  – Borrow snapshot space as needed

  You can modify the snapshot space recovery policy in **Group Manager → Group Configuration** , under the Default tab.

  If a snapshot has active iSCSI connections, the group closes the connections before deleting the snapshot.

  NOTE: In some cases, you might want to preserve the data in a snapshot that could be at risk of deletion. To preserve the data in a snapshot, you can clone the snapshot.
- Snapshot space warning percentage

  Percentage of the snapshot reserve that, when reached by in-use snapshot reserve, results in an event message. The default is 90 percent of the snapshot reserve. For example, if snapshot reserve space is 200MB and the warning level is 90 percent, a warning occurs when in-use snapshot reserve equals or exceeds 180MB.

Modify Snapshot Space Reserve Settings

If you change the snapshot space reserve, the values in the Pool Space table change. If the new snapshot space reserve value exceeds the capacity of the pool, the free pool space cell displays a negative value.

To modify the snapshot space reserve settings for a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Modify settings**.

The **Modify volume settings** dialog box opens.

4. In the dialog box, click **Space**. In the Snapshot Space section, modify the following values as needed:
   - Snapshot reserve – The reserve is the amount of space allocated to snapshot storage for the volume. It is expressed as a percentage of the volume's size. For example, if the volume size is 100GB and the snapshot reserve is 100 (the default), an additional 100GB of space will be allocated for snapshot storage. Changes to the amount of snapshot reserve space are immediately reflected in the Snapshot Reserve bar below, which shows the amount of free and used space.
   - Snapshot space warning percentage – The sliding pointer on the Snapshot Reserve bar determines how much snapshot reserve space will be consumed before a warning is issued that the volume is running out of snapshot reserve space. Click and drag the pointer to slide it to a new position. The default value is 90 percent.
5. Click **OK**.

NOTE: To change the groupwide default values for both the snapshot space reserve and the warning percentage:

1. Click **Group** → **Group Configuration**.
2. Click the **Defaults** tab.
3. In the **Snapshot space reserve** section of the Volume Settings panel, modify the settings as needed.

Settings specified here apply to newly created volumes only, not to existing volumes.

Modify Snapshot Space Recovery Policy

To modify the snapshot space recovery policy:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, below Snapshots click **Modify snapshot policy**.

   The **Modify snapshot policy** dialog box opens.
4. In the dialog box, select either **Set volume offline** or **Delete oldest snapshot**.

   If you select **Delete oldest snapshot**, you also have the option of selecting **Borrow snapshot space as needed**.
5. Click **OK**.

NOTE: To change the groupwide default values for the snapshot space recovery policy:

1. Click **Group** → **Group Configuration**.
2. Click the **Defaults** tab.
3. In the **Snapshot space recovery** section of the Volume Settings panel, modify the settings as needed.

Settings specified here apply to newly created volumes only, not to existing volumes.

Create a Snapshot

You can create a snapshot of a single volume at the current time. Snapshot creation occurs immediately, with no impact on volume availability or performance.

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Create snapshot**. The Create Snapshot dialog box opens.
4. (Optional) Type a description for the snapshot, up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. With the **Set snapshot online** checkbox, select whether to keep the snapshot offline (default) or set the snapshot online.
6. With the **Make snapshot read-write** checkbox, select whether to make the snapshot permission read-write (the default) or read-only.
7. Click **OK**.

The snapshot appears in the far-left panel, under the volume name.

The default snapshot name is the volume name followed by the date and time when you created the snapshot (for example, `dbasevolume-2014-03-25-15:31:14.7668`). Snapshots appear under a volume in the far-left panel listed by timestamp.

When you select a snapshot timestamp, its full name (volume and timestamp) appears in the GUI main window and in the Snapshot iSCSI Settings panel. You can modify the snapshot name. The new name can contain up to 127 characters. (Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.)

For information about the number of snapshots that can be created, see the *Dell EqualLogic PS Series Storage Arrays Release Notes*. When the maximum number of snapshots for a volume or group is reached, the system displays an error message and no more snapshots can be created.

Set a Snapshot Online or Offline

By default, a snapshot is offline. You can set a snapshot online, making it accessible to iSCSI initiators that match one of the snapshot’s access control policies.

If you set a snapshot offline, any current iSCSI connections to the snapshot are lost.

1. Click **Volumes**.
2. Expand **Volumes**.
3. Expand the volume name and then select the snapshot timestamp.
4. In the Activities panel, click **Set snapshot online** or **Set snapshot offline**.
5. Click **Yes** to confirm the choice to put the snapshot online or offline.

Clone a Snapshot to Create a New Volume

Cloning a snapshot creates a new standard volume, template volume, or thin clone volume. The new volume has a new name and new iSCSI target name, but the same reported size, pool, and contents that the original volume had at the time that you created the snapshot.

The group allocates space equal to the volume reserve that you specify for the new volume. If you reserve snapshot space for the new volume, the group allocates additional space.

The original snapshot still exists after the clone operation. The cloning process is accomplished through a wizard interface and is similar to creating a new volume, except that some attributes are predetermined by the source volume.

To clone a snapshot:

1. Click **Volumes**.
2. Expand **Volumes**.
3. Expand the volume name and then select a snapshot timestamp.
4. In the Activities panel, click **Clone snapshot**. The Clone Snapshot wizard opens.
5. Provide the required information in each step of the wizard, and click **Next**.
6. When you reach the final step, click **Finish**.

Modify a Snapshot Name or Description

The default snapshot name is based on the volume name and the time that you created the snapshot. You can modify this name and the optional snapshot description.

The following considerations apply:

- If you modify a snapshot name, and the alias (public name) is set to be the same as the snapshot name, the alias changes to match the new name.
- The iSCSI target name for the snapshot does not change if you change the snapshot name.

To modify a snapshot name or description:

1. Click **Volumes**.
2. Expand **Volumes** and then expand the volume name.
3. Select the snapshot timestamp.
4. In the Activities panel, click **Modify snapshot properties** to open the dialog box.
5. In the **General** tab, type the new snapshot name and, optionally, a description.
6. Click **OK**.

NOTE: Snapshot names can be up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

Modify the Snapshot Alias

You can modify the public alias (public name) for a snapshot. Some iSCSI initiators show the alias along with the iSCSI target name. The alias can help you identify the snapshot.

If the alias is used with an iSCSI connection, the alias name follows the volume name.

1. Click **Volumes**.
2. Expand **Volumes** and then expand the volume name.
3. Select the snapshot timestamp.
4. In the Activities panel, click **Modify snapshot properties** to open the dialog box.
5. Click the **iSCSI** tab.
6. Type the new alias in the **Public alias** field. Alias names can be up to 54 ASCII characters.
7. Click **OK**.

Modify Snapshot Permission

A snapshot can have read-only or read-write permission.

The following considerations apply:

- If you write to a snapshot, it might no longer represent the contents of the base volume at the time of snapshot creation.
- To change the permission of an online snapshot to read-only, first set the snapshot offline.

To modify the permission for a snapshot:

1. Click **Volumes**.
2. Expand **Volumes** and then expand the volume name.
3. Select the snapshot timestamp.
4. In the Activities panel, click **Set access type**. The Select Access Type dialog box opens.
5. Select the permission for the snapshot (either **Set read-write** or **Set read-only**).
6. If your environment supports multiple initiators accessing a volume, you can check the **Allow simultaneous connections from initiators with different IQNs** checkbox. This option is disabled by default.
7. Click **OK**.

Delete Snapshots

You can delete unwanted snapshots. If you delete a snapshot that is part of a snapshot collection or a custom snapshot collection, the collection's Integrity status changes to `incomplete`.

NOTE: You must take a snapshot offline before you can delete it. To take the snapshot offline, click the snapshot’s timestamp and click Set snapshot offline in the Activities panel.

1. Click **Volumes**.
2. Expand **Volumes**.
3. Expand the volume name and select the snapshot timestamp that you want to delete.

4. In the Activities panel, click **Delete snapshot**.
5. When prompted to confirm the deletion, click **Yes**.

Restore a Volume from a Snapshot

You can restore a volume from a snapshot, and replace the data in the current volume with the volume data at the time you created the snapshot. The snapshot still exists after the restore operation.

The following considerations and constraints apply when restoring a volume from a snapshot:

- All members that contain data from a volume or snapshot must be online.
- A template volume cannot be restored from a snapshot.
- Before the restore operation starts, the group automatically creates a snapshot of the current volume.
- A synchronous replication (SyncRep) volume cannot be restored from a snapshot if the snapshot's size is different from that of the volume.

To restore a volume from a snapshot:

1. Disconnect any iSCSI initiators from the volume. Follow the instructions for your operating system and initiator.
2. Click **Volumes**.
3. Expand **Volumes** and then select the volume name.
4. Click **Set offline** and confirm that you want to set the volume offline.
5. Click **Restore volume** to open the Restore Volume from Snapshot dialog box.
6. Select the timestamp of the snapshot that you want to restore and click **OK**. The Restore Volume from Snapshot dialog box opens.
7. If you want the volume to be put online after it is restored, select the **Set volume online after restore is complete** checkbox.
8. Click **Yes** to initiate the restore process. If you selected the **Set volume online...** checkbox in step 7, the volume will come online when the restore procedure completes.

The restored volume has the same name and iSCSI target name as the original volume.

About Snapshots and NAS Container Data

Snapshots are read-only, point-in-time copies of NAS container data. Snapshots use a redirect-on-write method to track NAS container changes. Deleting a NAS container deletes any snapshots of the NAS container.

Before you can create a snapshot, you must specify the amount of snapshot reserve for the NAS container. Snapshot reserve is specified as a percentage of the NAS container space. Therefore, user data and snapshots compete for the same NAS container space.

NOTE: Select the Enable data reduction checkbox to activate the Modify policy button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.

Snapshot creation is instantaneous. You can schedule snapshots to create them regularly.

Snapshot reserve is used only if NAS container changes occur. When free snapshot reserve is consumed, the NAS cluster deletes the oldest snapshot before creating a new one.

Clients can easily retrieve files in a snapshot, without administrator intervention. If necessary, you can restore a NAS container from a snapshot.

Create a NAS Container Snapshot

To create a snapshot of a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Create snapshot**.

3. In the **Name** field, type a snapshot name:
   - A snapshot name can contain up to 229 characters, including letters, numbers, and underscores. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - If you do not assign a snapshot name, the NAS cluster generates a name automatically, based on the NAS container name and the timestamp.
4. Click **OK**.

The snapshot appears (identified by its timestamp) when you expand the NAS container name in the far-left panel.

About Deleting NAS Container Snapshots

You can manually delete a snapshot when it is no longer needed.

The NAS cluster automatically deletes one or more of the oldest snapshots for a NAS container when:

- The number of snapshots created from a schedule exceeds the max-keep limit.
- A NAS container data deletion or write causes the in-use snapshot space to be larger than the size of the snapshot reserve.

  Select the **Enable data reduction** checkbox to activate the **Modify policy** button. Enabling data reduction permanently removes the snapshot reserve functionality from the NAS container.

In addition, the NAS cluster also automatically deletes NAS container snapshots when:

- You delete a NAS container. All the snapshots configured for that NAS container are also deleted.
- You restore a NAS container from a snapshot. All the snapshots created after the snapshot that was used to restore the NAS container are deleted.

Delete a NAS Container Snapshot

To delete a NAS container snapshot:

1. Click **NAS** and expand **NAS Cluster**.
2. Expand **Local Containers**, then select and expand a NAS container name.
3. Select the snapshot that you want to delete.
4. In the Activities panel, click **Delete snapshot**.
5. Click **Yes** to confirm that you want to delete the snapshot.

Restore a NAS Container from a Snapshot

To restore a NAS container from a snapshot:

1. Click **NAS** and expand **NAS Cluster**.
2. Expand **Local Containers** and select a NAS container.
3. Click **Restore NAS container**.
4. Select the snapshot to restore and click **OK**.
5. Click **Yes** to confirm that you want to restore the NAS container from the snapshot.

After the restore operation completes, the NAS cluster deletes any snapshots that were created after the snapshot that was used to restore the NAS container. Snapshots that were created before the snapshot that was used for the restoration are not affected.

NOTE: When you restore a NAS container from a snapshot, any currently mounted NFS exports will produce a `Stale NFS file handle` error message. You must unmount and then remount the NFS exports.

Retrieve a File from a NAS Container Snapshot

1. Locate the .snapshot/[*snapshot_name*] directory in the affected NFS export or SMB share.
2. Locate the file in the snapshot directory.
3. Copy the file to the new location.

Modify the Name of a NAS Container Snapshot

When you create a snapshot, the name automatically assigned to the snapshot is based on the volume name, and includes a timestamp and an identification number.

To modify the name of a NAS container snapshot:

1. Click **NAS** and expand **NAS Cluster**.
2. Expand **Local containers** and select the container that is associated with the snapshot that you want to modify.

   A plus sign appears next to the container names that are associated with one or more snapshots.
3. Expand the container and select the snapshot to modify.
4. In the Activities panel, click **Modify snapshot properties** to open the dialog box.
5. In **Snapshot name**, modify the snapshot name.

   A snapshot name can contain up to 229 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
6. Click **OK**.

About Snapshot Collections

A standard snapshot collection results when you create a snapshot of a volume collection. This process creates a snapshot of each volume in the collection. The ability to create snapshots of multiple volumes simultaneously is useful when you want to protect data in multiple, related volumes. As with single-volume snapshots, you can create a schedule for a standard snapshot collection.

An alternative method for creating snapshots of multiple volumes without using a volume collection is to create a custom snapshot collection. Custom snapshot collections can include any volume in the group, so you manually select which volumes you want to include when you configure the collection. Unlike the standard snapshot collection, you cannot create a schedule for custom snapshot collections.

Create a Snapshot Collection

The following prerequisites and considerations apply before you can create a snapshot collection:

- You need to have at least one volume collection.
- You must first allocate snapshot reserve for each volume in the volume collection.

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection name.
3. In the Activities panel, click **Create snapshot collection**. The Create Snapshot Collection dialog box opens.
4. (Optional) Type a description of the collection, up to 127 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
5. Click **OK**.

   In a few moments, a new snapshot collection will appear under the volume collection name.

Create a Custom Snapshot Collection

You can create snapshots of multiple volumes in a single operation, without using a volume collection. The set of snapshots, one for each volume, is called a custom snapshot collection.

Prerequisites

- Allocate snapshot reserve for each volume in the volume collection.
- You need the following information:
  - Volume name of each volume that you intend to include in the collection.
  - Collection name, up to 127 characters, with no spaces.

  Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

- (Optional) Collection description, up to 127 characters, including colons, hyphens, and periods.

  Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

Procedure

To create a custom snapshot collection:

1. Click **Volumes**.
2. Click **Custom Snapshot Collections** in the Volumes list.
3. In the Activities panel, click **Create custom collection** to open the dialog box.
4. Provide the requested information in each step of the wizard and click **Next**.
5. Click **Finish** to create the collection or click **Back** to make changes.

The default custom snapshot collection name is the volume collection name followed by the date and time when you created the snapshots (for example, UserData-2014-05-14-15:08:18.3).

Modify a Snapshot Collection Name or Description

Determine the following information before making any changes:

- Whether you are modifying a snapshot collection or a custom snapshot collection
- Timestamp of the collection

1. Click **Volumes**, then either:
   - Expand **Volume Collections**, then expand the collection and select the timestamp for the snapshot collection.
   - Expand **Custom Snapshot Collections** and select the timestamp for the custom snapshot collection.
2. In the Activities panel, click either **Modify snapshot collection** or **Modify custom collection**.
3. In the dialog box, modify the snapshot collection name (required) and description (optional):
   - A snapshot collection name can be up to 127 characters, with no spaces.

     Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - A snapshot collection description can be up to 127 characters, including colons, hyphens, and periods.

     Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
4. Click **OK**.

Delete a Snapshot Collection

Deleting a snapshot collection also deletes the snapshots in the collection. Before deleting a snapshot collection, determine the following information:

- Whether you are deleting a snapshot collection or a custom snapshot collection
- Timestamp of the collection (to make sure you delete the correct one)

1. Click **Volumes**, then either:
   - Expand **Volume Collections** and select the timestamp for the snapshot collection that you want to delete.
   - Expand **Custom Snapshot Collections** and select the timestamp for the custom snapshot collection that you want to delete.
2. In the Activities panel, click **Delete snapshot collection** or **Delete custom collection**.
3. When prompted, confirm that you want to delete the snapshot collection or the custom snapshot collection by clicking **Yes**.

About Snapshot Space Borrowing

Snapshot space borrowing enables the system to temporarily increase the available snapshot space for a volume by borrowing from other sources.

If borrowed space is needed for other functions, the firmware might delete snapshots that are borrowing space. Because of this potential deletion, you should always be aware of when snapshots are borrowing space.

NOTE: Borrowed space is intended to help during peaks of activity when more space is needed temporarily. It does not take the place of carefully provisioning snapshot reserves.

A regular (standard) volume has a fixed volume reserve; this volume reserve for regular volumes does not change when its usage increases. With a thin-provisioned volume, the system automatically increases volume reserve as more data is written to that volume. As the volume reserve increases, so does the allotted snapshot space for that volume.

Snapshot borrowing allows a thin-provisioned volume to keep its snapshots when snapshot reserve shrinks due to volume unmapping. When a volume is unmapped, the volume reserve shrinks; because snapshot reserve is a percentage of volume reserve, snapshot reserve also shrinks.

Activate Snapshot Space Borrowing

NOTE:

- **If you are using snapshot space borrowing on a new group in which all the members are running the latest version of the Group Manager software, snapshot borrowing is available immediately.**
- **If you are upgrading the members in the group from a previous version of the firmware, the Disallow Downgrades function must be set on each member in the group before snapshot borrowing can be activated. This function is set automatically when all members in the group are running Group Manager firmware version 6.0 or later.**
- **All members in the group must be running Group Manager firmware version 6.0 or later. Snapshot space borrowing does not work in environments running mixed versions of the firmware.**
- **You can set a system default policy for new volumes. For existing volumes, you must configure the settings on each individual volume.**

To set a system default policy for snapshot borrowing for new volumes:

1. Click **Group → Group Configuration**.
2. Click the **Defaults** tab.
3. In the **Snapshot space reserve** section, specify a value (in percent) for the snapshot reserve (the default value is 100 percent, but it can exceed 100).
4. In the Volume Settings panel, navigate to the **Snapshot space recovery** section.
5. Select **Delete oldest snapshot**, if it has not already been selected. This option enables snapshot space borrowing; it will not work without this option selected.
6. Make sure that the **Borrow snapshot space as needed** checkbox is selected.

When snapshot borrowing is enabled, the snapshot in-use space is allowed to exceed the snapshot reserve space. Some warning messages about snapshot reserve in-use space on the volume are suppressed when snapshot space borrowing is in effect; however, warning messages are still issued when in-use snapshot space exceeds the warning threshold.

NOTE: Snapshot space borrowing cannot be enabled if the Set volume offline option is selected under Snapshot space recovery. If you disable the system default policy for snapshot space borrowing for a newly created volume, existing volumes will still have snapshot space borrowing enabled.

Control Snapshot Space Borrowing

You can control whether or not a volume is allowed to borrow space for snapshots. Snapshot space borrowing enables you to temporarily increase the available snapshot space for a volume by borrowing space from other sources. Borrowing can help prevent the oldest snapshots from potentially being deleted when the volume's allocated snapshot reserve is depleted.

To control snapshot space borrowing for existing volumes:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Snapshots section of the Activities panel, click **Modify snapshot policy**.
4. In the dialog box, select either **Set volume offline** or **Delete oldest snapshot**.

   NOTE: Snapshot space borrowing cannot be enabled if the Set volume offline option is selected. The Delete oldest snapshot option enables snapshot space borrowing.

   NOTE: You cannot select Set volume offline for recovery volumes.
5. To enable snapshot space borrowing, select **Borrow snapshot space as needed**.

   To disable snapshot space borrowing, clear the checkbox.

   NOTE: You can also control snapshot space borrowing under the Group Configuration Defaults tab. Settings specified here apply to newly created volumes only, not to existing volumes.

About Replication

The replication technology provided by PS Series firmware enables you to copy volume data from one group to another, protecting the data from a variety of failures, ranging from the destruction of a volume to a complete site disaster, with no effect on data availability or performance.

Volume replication between different groups provides protection against data loss. If a volume is destroyed, you can fail over to the recovery group and recover data from a replica. Users can then resume access to the recovery volume. After the original volume becomes available, you can fail back to the original group.

Traditional replication is different from synchronous replication (SyncRep). You cannot enable synchronous replication on a volume for which traditional replication is configured, and you cannot enable traditional replication on a volume for which synchronous replication is configured.

NOTE: A group can have multiple replication partners. However, you can replicate a volume to only one replication partner at a time. Choose the replication configuration that is right for your environment.

You can use PS Series replication functionality alone or with Auto-Snapshot Manager (for Linux and Windows), Dell Virtual Storage Manager for VMware (formerly ASM/VE), or Storage Replication Adapter for VMware Site Recovery Manager. See the related product documentation for details.

The first step in implementing replication is to configure the replication partners. Replication partners can consist of two or more groups. In the simplest case, a primary group contains the volume you want to replicate and a secondary group stores the replicated data. Each partner plays a role in the replication of a volume, and you can monitor replication activity from either partner:

- Primary group

  Location of the volume — The primary group administrator configures the secondary group as a replication partner and initiates the volume replication operation. Replication of the volume is considered outbound from the view of the primary group.
- Secondary group

  Location of the volume’s replica set — The secondary group administrator configures the primary group as a replication partner and delegates space for storing replicas from the primary group. Replication of a volume is considered inbound from the view of the secondary group (sometimes called the destination group).

Mutual authentication using passwords provides security between partners.

After you configure the two groups as replication partners, you can configure a volume or volume collection for replication, specifying the replication partner, the local group space for the replication operation, and the remote partner space for storing the replicas.

The first time you replicate a volume, the primary group copies the entire contents of the volume to the secondary group. For subsequent replication operations, the primary group copies only the data that changed since the previous replication operation started.

Eventually, the oldest replicas are deleted from the replica set to free space for new replicas. The amount of space that you allocate for storing replicas limits the number of replicas you can keep on the secondary group.

NOTE: To ensure that a complete copy of volume data exists on the secondary group, the most recent, complete replica of a volume cannot be deleted.

To access or recover volume data from replicas, you can:

- Clone an individual replica to create a new volume on the secondary group.
- Promote the replica set to a recovery volume (and snapshots) on the secondary group and configure initiators to connect to the recovery volume.

  If the primary group becomes available, you can replicate the recovery volume to the primary group and then fail back to the primary group, returning to the original configuration.

NOTE: Replication is used primarily for disaster recovery and does not take the place of a comprehensive backup strategy.

You can manually delete replicas and replica sets that are no longer needed. You cannot delete the most recent, complete replica from a replica set, but you can delete the entire replica set, which disables replication on the volume.

NOTE: For information about NAS replication options, refer to NAS Volume Replication Operations.

Replication Types

PS Series firmware provides the following methods for automatically replicating block volumes to provide protection against accidental data loss:

- Traditional replication (often referred to as Replication or Auto-Replication) is a point-in-time process that is conducted between two groups, often in geographically diverse locations. Replication provides protection against a regional disaster such as an earthquake or hurricane.
- Synchronous replication (also known as SyncRep) is a real-time process that simultaneously writes volume data across two different pools within the same PS Series group. This method is useful for maintaining two copies of a volume's data in the same data center, or dispersed to two different facilities on the same campus or in the same metropolitan area.

  NOTE: You cannot enable synchronous replication on a volume for which traditional replication is configured, and you cannot enable traditional replication on a volume for which synchronous replication is configured.

In addition to these two replication methods, the Dell Fluid File System (FluidFS) release available on FS7500 and FS76x0 series NAS appliances allows you to replicate NAS containers between EqualLogic NAS clusters. This point-in-time NAS replication uses space-efficient snapshots of NAS containers to replicate file data. You can manage both block and NAS replication using the Group Manager interface.

NOTE: Throughout the firmware documentation, generic references to replication always refer to traditional replication.

A third type of replication, called manual transfer replication, accommodates less common instances where the network link between groups might be insufficient to support large data transfers. This method is typically done as a precursor to enacting traditional replication, so the initial full-volume transfers do not have to be done over the network. Manual transfer replication involves copying the data to external media, then physically transferring it and recopying it to the other member. See the *Dell EqualLogic Manual Transfer Utility Installation and User's Guide* for more information.

Plan Your Volume Replication Environment

To help ensure successful replication, for each volume that you want to replicate follow these steps to set up your replication environment:

1. Gather the following information to help you determine how much replication space you need:
   - Number of replicas you want to keep
   - Average time span between each consecutive replica

- Reported size of the volume
- Whether thin-provisioned or not
- Estimated rate of volume changes (depends on volume usage)

2. Make sure that the primary group has enough free pool space for the local replication reserve for each replicated volume.
3. Identify a replication partner (secondary group) to store the volume replicas. This secondary group must meet the space and network connectivity requirements.
4. Verify on the primary group that the secondary group is configured as a replication partner.
5. Verify on the secondary group:

   - The primary group is configured as a replication partner.
   - The correct amount of space is delegated to the primary group for storing replicas of primary group volumes. Increase that space if necessary.

6. On the primary group, configure the volume for replication, specifying the appropriate replication space values, and verifying the replication partnership.
7. On the primary group, replicate the volume.
8. (Optional) Set up a schedule to create replicas on a regular basis.
9. Monitor each replication operation and make sure it is successful.

   If the replication operation is not successful, identify and correct the problem. For example, you might need to increase network bandwidth or increase replication space.

10. Monitor the number of replicas and replication space usage over time. The goal is to keep a specific number of replicas without wasting replication space.

    - If replicas are deleted before you reach the number of replicas you want to keep, you might want to increase the replica reserve percentage for the volume.
    - If you are keeping an excessive number of replicas, you might want to decrease the replica reserve percentage for the volume.

Replication to One Partner

One replication partner replicates volumes to another partner. For example, in Figure 15. Replication to One Partner, GroupA replicates Volume1 and Volume2 to GroupB. GroupA is the primary group, and GroupB is the secondary group.

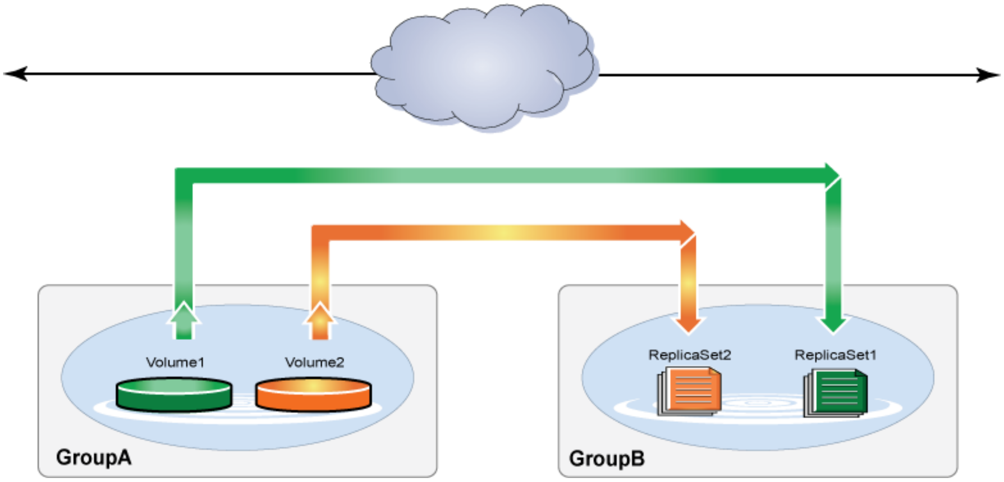

Figure 15. Replication to One Partner

Replication to Multiple Partners

One replication partner replicates different volumes to different partners. For example, in Figure 16. Replication to Multiple Partners, GroupA replicates Volume1 to GroupC, and GroupA replicates Volume2 to GroupB. GroupA is the primary group, and GroupB and GroupC are secondary groups.

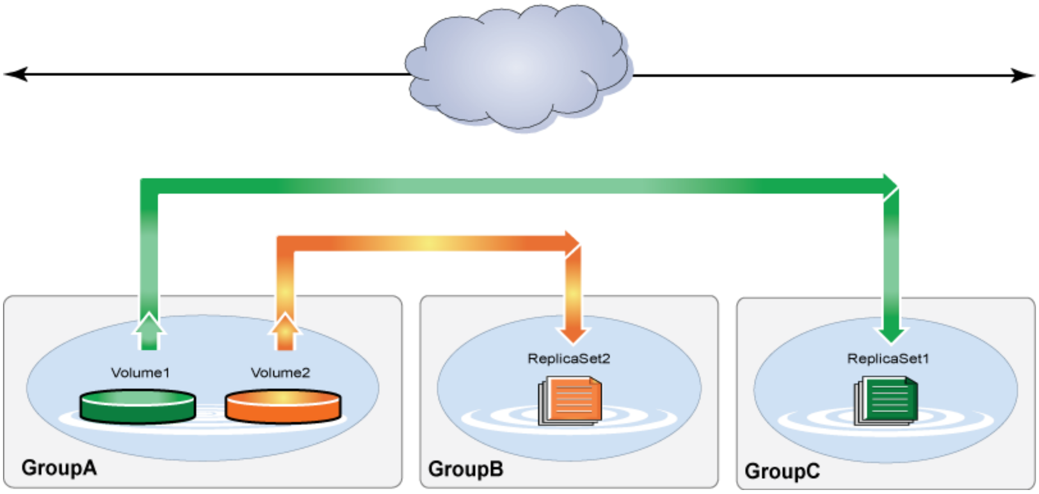

Figure 16. Replication to Multiple Partners

Reciprocal Replication Between Partners

Both partners replicate volumes to each other. For example, in Figure 17. Reciprocal Replication Between Partners, GroupA replicates Volume1 to GroupB, and GroupB replicates Volume2 to GroupA. For the replication of Volume1, GroupA is the primary group and GroupB is the secondary group. For the replication of Volume2, GroupB is the primary group and GroupA is the secondary group.

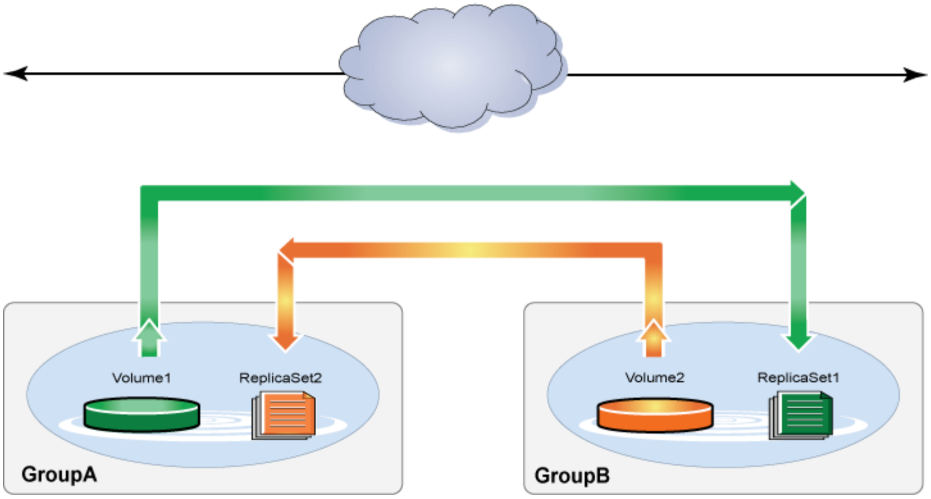

Figure 17. Reciprocal Replication Between Partners

Centralized Replication

Multiple partners replicate volumes to another partner. For example, in Figure 18. Centralized Replication, GroupA and GroupB replicate volumes to Group C. In this configuration, GroupA and GroupB are primary groups, and GroupC is the secondary group.

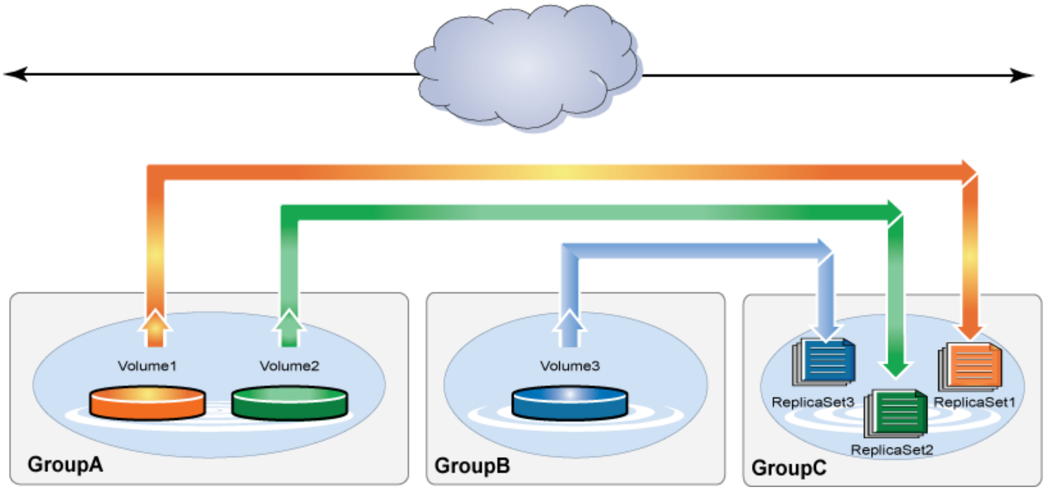

Figure 18. Centralized Replication

About Replication Space

Volume replication between partners requires space on both the primary group (the volume location) and the secondary group (the replica location). These space requirements are classified in the following way:

- Local replication reserve

  Each volume requires primary group space for use during replication and, optionally, for storing the failback snapshot.
- Delegated space

  To provide space for storing replicas, the secondary group delegates space to the primary group. All primary group volumes that you replicate to the secondary group share the delegated space.

  Each volume is assigned a portion of delegated space, called the replica reserve. The replica reserve for a volume limits the number of replicas you can keep. When replica reserve is consumed, the oldest replicas are deleted to free space for new replicas.

To make sure replication operations complete and to keep the desired number of volume replicas, you must allocate sufficient replication space.

To determine the optimal amount of replication space, Dell recommends that you set up replication using the default space values, monitor activity over some time period, analyze the space usage, and make adjustments. These precautions will help you keep the desired number of replicas while using replication space efficiently.

About Local Replication Reserve

Each replicated volume requires primary group space, called local replication reserve. See Figure 19. Local Replication Reserve.

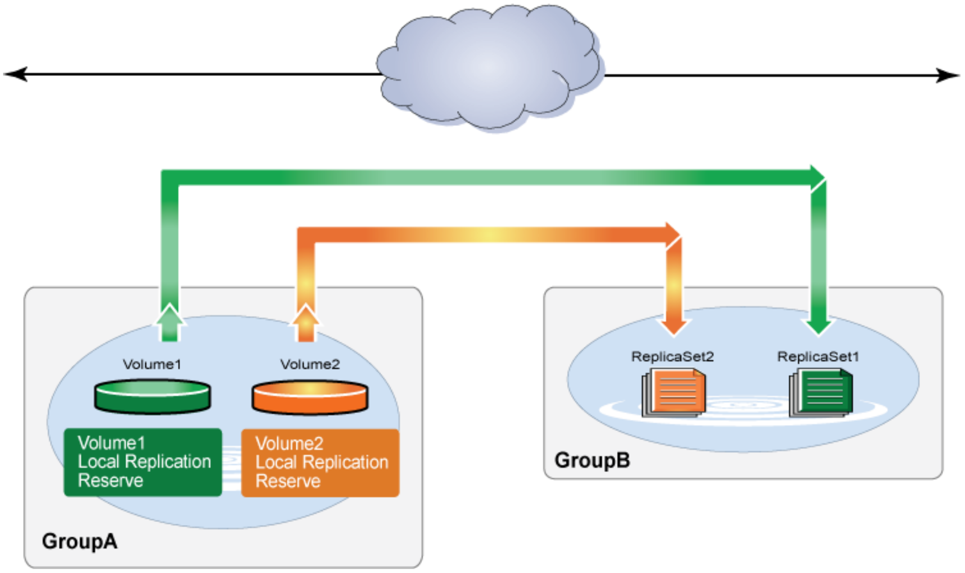
Volume1
Volume2
Volume1 Local Replication Reserve
Volume2 Local Replication Reserve
GroupA
ReplicaSet2
ReplicaSet1
GroupB

Figure 19. Local Replication Reserve

Local replication reserve has two purposes:

- Preserve the contents of the volume at the time replication started. The primary group creates a snapshot of the volume in the local replication reserve to preserve the contents of the volume at the time replication started. If the volume changes occur during replication, the snapshot tracks those changes, consuming more local replication reserve.

  When replication completes, the primary group deletes the snapshot, freeing the space, unless you chose the option to keep the failback snapshot.
- Store the failback snapshot (optional). The failback snapshot for a volume can expedite volume failback operations.

  If you choose to keep the failback snapshot when configuring a volume for replication, the primary group does not delete the snapshot in the local replication reserve when replication completes. Instead, it becomes the failback snapshot. As volume changes occur between replication operations, the failback snapshot consumes more local replication reserve.

  After each replication completes, the primary group replaces the failback snapshot to update the failback baseline. Therefore, the volume data represented by the failback snapshot on the primary group always matches the volume data represented by the most recent, complete replica on the secondary group.

  If you fail over to the secondary group and write to the recovery volume, you can fail back to the primary group by replicating only the changes made to the recovery volume if the failback snapshot still exists. If the failback snapshot does not exist, you must replicate the entire volume contents to the primary group to complete the failback operation.

It is important to allocate sufficient local replication reserve to ensure that replication operations complete and, optionally, to maintain the failback snapshot.

If not enough free local replication reserve is available to complete a replication operation, one of the following situations occurs:

- If you enabled the option to borrow free pool space, and sufficient free pool space is available (at least 100GB per member in the pool), replication continues. The primary group generates an informational message, specifying that it is temporarily using free pool space during the replication.
- If you did not enable the option to borrow free pool space (or if you enabled the option, but not enough free pool space is available), the primary group cancels the replication and generates an event message, stating that the replication was canceled.

If not enough free local replication reserve is available to maintain the failback snapshot, one of the following situation occurs:

- If you enabled the option to borrow free pool space, and free pool space is available, the primary group generates an informational message specifying that it is temporarily using free pool space.

- If you did not enable the option to borrow free pool space (or if you enabled the option, but not enough free pool space is available), the primary group deletes the failback snapshot and generates an event message. To reestablish the failback snapshot, increase the local replication reserve and replicate the volume.

In addition, if you attempt to replicate a recovery volume to the primary group, and not enough local replication reserve is available to store the changes, the primary group generates an event message advising you to increase the space.

Guidelines for Sizing Local Replication Reserve for Use During Replication

To size the portion of local replication reserve for use during replication, follow these guidelines:

- Recommended value

  A value of 100 percent ensures sufficient local replication reserve even if the entire volume changes during a replication operation.
- Space-efficient value

  If few volume changes are expected during an average replication operation, use a value less than 100 percent.

To obtain an appropriate value, estimate the average volume change that occurs during a typical replication operation. Then, use the following equation, where 5 percent is the minimum local replication reserve:

```
5 percent + change rate
```

For example, if you estimate that at most 10 percent of the volume changes, a value of 15 percent might be appropriate (5 percent plus 10 percent).

If you use a local replication reserve value that is less than 100 percent, sufficient space might not be available to complete a replication operation (for example, if more volume changes than expected occur during a replication operation). Therefore, Dell recommends that you select the option that allows you to borrow free pool space if not enough local replication reserve is available to complete a replication.

Guidelines for Sizing the Local Replication Reserve for the Failback Snapshot

To size the portion of local replication reserve used to maintain the failback snapshot, follow these guidelines:

- Recommended value

  A value of 100 percent ensures sufficient local replication reserve even if the entire volume changes between replication operations.
- Space-efficient value

  If, on average, few volume changes are expected between replication operations, use a value less than 100 percent.

To obtain an appropriate value, estimate the average volume change that occurs between consecutive replication operations. Then, use the following equation, where 5 percent is the minimum local replication reserve:

```
5 percent + change rate
```

For example, if you estimate that at most 20 percent of the volume changes, a value of 25 percent might be appropriate (5 percent plus 20 percent).

If you use a local replication reserve value that is less than 100 percent, sufficient space might not be available to maintain the failback snapshot. Therefore, Dell recommends that you select the option that allows you to borrow free pool space if not enough local replication reserve is available.

Guidelines for Sizing Local Replication Reserve

On the primary group, you can specify the value of the local replication reserve and whether to keep the failback snapshot when you configure a volume for replication. You can also enable the option that allows you to borrow free pool space if not enough local replication reserve is available. You can later modify these settings as needed.

The local replication reserve size is based on a percentage (5 percent to 200 percent) of the volume reserve. For a thin-provisioned volume, the volume reserve size changes dynamically based on volume usage; therefore, the local replication reserve size also changes.

The recommended local replication reserve percentage depends on whether you want to keep the failback snapshot:

- No failback snapshot

  Specify 100 percent for the local replication reserve.
- Keep the failback snapshot

  Specify 200 percent for the local replication reserve.

Using the recommended values might not be the most efficient use of local replication reserve. Ideally, you want to allocate only enough space to meet the volume requirements. However, specifying too little space can prevent successful replication.

The optimal value for local replication reserve depends on the volume change rate, the replication frequency, and whether you are keeping the failback snapshot. The volume change rate can be difficult to estimate.

If you want to use *less* than the recommended value for local replication reserve, follow these guidelines:

- No failback snapshot

  Size the local replication reserve based only on its use during replication.
- Keep the failback snapshot

  Size the local replication reserve based on its use during replication and also for maintaining the failback snapshot. Then, combine the two values.

About Delegated Space and Replica Reserve

Replicas are stored in space that the secondary group delegates to the primary group.

For example, if you want to replicate GroupA volumes and GroupB volumes to GroupC, GroupC must delegate space to GroupA and GroupB. See Figure 20. Delegated Space.

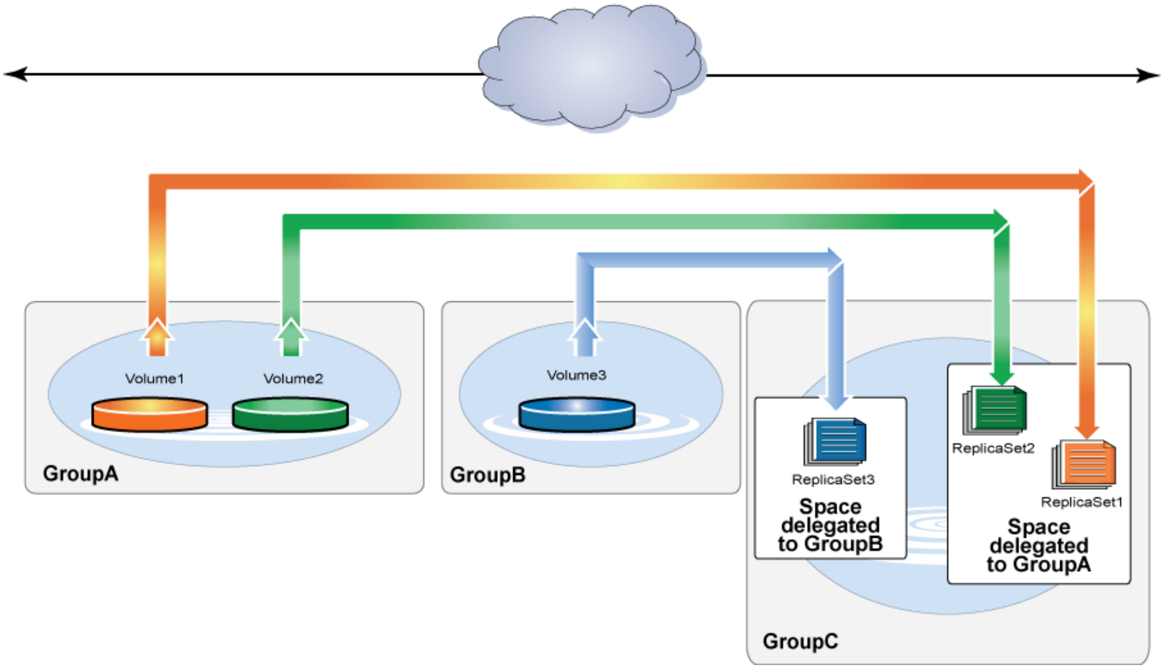

Figure 20. Delegated Space

The secondary group administrator delegates space to the primary group when configuring the group as a replication partner. The administrator can modify the partner configuration and increase or decrease delegated space as needed.

When the primary group administrator configures a volume for replication, a portion of that delegated space is reserved for the volume. This space, called replica reserve, limits the number of replicas for that volume that you can keep on the secondary group. You can modify the volume replication configuration and increase or decrease the replica reserve value as needed. (Note that replication borrowing can also affect the size of the reserve. See About Replication Borrowing for more information.)

To determine the correct amount of space that the secondary group needs to delegate to the primary group, you must obtain the replica reserve requirement for each primary group volume that you are replicating to the secondary group. Replica reserve is based on a percentage of the volume's volume reserve. For more information, see Guidelines for Sizing Delegated Space.

Manage Space Delegated to a Partner

You can modify the space delegated to a partner, subject to the following restrictions:

- You cannot decrease the space delegated to a lower capacity than is currently reserved for the partner's replicas.
- The Pool Space table shows how pool space is currently used and how much space is used after the change.

  If more than one pool is configured to have delegated space, the first row of the table shows the overall statistics for delegated space usage across all pools. The lower rows show this information for the currently selected pool.

  If the new delegated space exceeds the capacity of a pool, the color of the table cell showing free pool space changes to red.

For additional information, see Guidelines for Sizing Delegated Space.

To modify the space delegated to a partner:

1. Click **Replication** and expand the replication partner.
2. Click **Volume Replication**.
3. In the Activities panel, click **Manage delegated space** to open the Manage Delegated Space dialog box.
4. Enter the new delegated space value. (If the group has more than one pool, you can specify a new value for each one as needed.)
5. Click **OK**.

About Replica Reserve Usage

To provide space for partner replicas, the secondary group delegates space to the partner.

When you configure a volume for replication, you assign a portion of delegated space to the volume. This space, called replica reserve, limits the number and size of volume replicas that you can keep on the secondary group.

It is important to correctly size the replica reserve for a volume. Too little replica reserve might prevent you from keeping the desired number of replicas. However, delegated space is limited, so keeping too many replicas might not be an efficient use of delegated space.

When you configure a volume for replication, you specify the replica reserve size as a percentage (minimum 105 percent) of the replica volume reserve, which approximates in-use volume space. As volume usage increases, the replica reserve also increases, providing more space for replicas.

For example, if you specify 200 percent for the replica reserve, and the replica volume reserve is 2.5GB, the replica reserve size is 5.0GB. If the replica volume reserve increases to 6.0GB, the replica reserve size increases to 12.0GB.

The replica volume reserve has a maximum size (the reported volume size for volumes that are not thin-provisioned, or the maximum in-use space value for thin-provisioned volumes); therefore, the replica reserve has a maximum size.

To properly size replica reserve, you must understand how replica reserve space is used.

Replica Reserve Usage – First Replication

1. The primary group determines how much volume data to replicate. For the first replication operation, the primary group must copy all the volume data.

2. The primary group increases the replica reserve if the replica volume usage increased since you enabled replication on the volume.

   NOTE: If delegated space is too small to hold the increased, replica reserve, the primary group generates an event message and replication pauses. Replication resumes automatically when delegated space is large enough to hold the reserve.

3. The primary group copies the contents of the volume to replica reserve, decreasing the amount of free replica reserve. For example, if the volume consists of 10GB of data, free replica reserve decreases by 10GB.

At this point, the replica reserve contains one replica, which is the most recent, complete replica.

Replica Reserve Usage – Subsequent Replications

1. The primary group determines how much volume data to replicate. For replication operations other than the first, the primary group copies only the data that changed since the previous complete replication (the deltas).
2. The primary group increases the replica reserve if the replica volume usage increased since the previous replication operation.

   NOTE: If delegated space is too small to hold the increased, replica reserve, the primary group generates an event message and replication pauses. Replication resumes automatically when delegated space is large enough to hold the reserve.

3. If not enough free replica reserve is available, the system tries to borrow space to fit the new replicas. If any objects are borrowing space, they might be deleted to make room for the new replicas. Also, borrowing can provide more space, but only temporarily. Replicas could still be deleted.

   The most recent, complete replica is never deleted automatically, ensuring that you always have a current copy of volume data on the secondary group.

4. The primary group copies the volume data to replica reserve, decreasing the amount of free replica reserve. For example, if the replication transferred 5GB of new data, free replica reserve decreases by 5GB.

To make sure that replication operations complete and keep the desired number of volume replicas, it is important to allocate sufficient delegated space and specify the correct replica reserve percentage. The amount of delegated space that you need depends on the replica reserve requirements for all the replicated volumes from a partner.

About Replica Volume Reserve

Each replicated volume has a replica volume reserve, which approximates the amount of in-use volume space. The value of the replica volume reserve is used to allocate replica reserve for a volume.

When you configure a volume for replication, the primary group sets the initial value of the replica volume reserve:

- For volumes that are not thin-provisioned, 10 percent of the reported volume size. For example, if you have a 10GB volume that is not thin-provisioned, the initial replica volume reserve is 1GB (10 percent of 10GB).
- For thin-provisioned volumes, the current volume reserve. For example, if you have a 10GB volume that is thin-provisioned with a volume reserve of 2.5GB, the initial replica volume reserve is 2.5GB.

The initial value of the replica volume reserve appears in the Configure Replication – General Settings dialog box.

At the start of each replication operation, the primary group determines whether to increase the value of the replica volume reserve:

- For volumes that are not thin-provisioned, 10 percent of the reported volume size.

  For example, if you have a 10GB volume that is not thin-provisioned, the initial replica volume reserve is 1GB (10 percent of 10GB).
- For thin-provisioned volumes, the current volume reserve.

  For example, if you have a 10GB volume that is thin-provisioned with a volume reserve of 2.5GB, the initial replica volume reserve is 2.5GB.

The size of the replica reserve for a volume is based on a percentage of the replica volume reserve. Therefore, as the replica volume reserve increases, the replica reserve also increases, providing more space for replicas up to a limit.

Guidelines for Sizing Replica Reserve for a Volume

To determine the amount of space that the secondary group must delegate to the primary group, you must obtain the replica reserve requirement for each primary group volume that you are replicating to the secondary group.

When you configure a volume for replication, you specify the replica reserve size as a percentage (minimum 105 percent) of the replica volume reserve, which approximates in-use volume space.

As volume changes occur, the replica volume reserve increases; therefore, the replica reserve increases, providing more free space for replicas up to a limit.

NOTE: Replica reserve can increase automatically or by administrator action only if free delegated space is available. See Guidelines for Sizing Delegated Space.

Ideally, you want to allocate only enough replica reserve to store the desired number of replicas. In general, the higher the replica reserve percentage, the more replicas you can store. However, specifying a high percentage for all volumes might not be the most efficient use of delegated space.

The optimal value for replica reserve depends on the volume change rate, which can be difficult to estimate, and the replication frequency. See How Volume Changes Affect Replication Space.

Use the following guidelines for sizing replica reserve:

- Recommended value

  Dell recommends that you specify 200 percent for the replica reserve. Specifying 200 percent guarantees that you can store at least two replicas, assuming enough delegated space is available for the replica reserve to reach its maximum size. If the replica reserve cannot reach its maximum size due to lack of delegated space, you are not guaranteed two replicas.

  If you want to guarantee more than two replicas, specify a higher percentage.
- Space-efficient value

  For volumes that are not frequently modified, you might be able to keep the desired number of replicas by using a replica reserve value that is less than 200 percent.

  To obtain an appropriate replica reserve value, estimate the average volume change that occurs between replication operations. Then, use the following calculation, where 105 percent is the minimum replica reserve value:

  ```
  105 percent + [change rate x (number of replicas to keep -1)]
  ```

  For example, if you estimate that at most 20 percent of the volume changes between replication operations, and you want to keep three replicas, specify 145 percent for the replica reserve value.

After replication begins, monitor the number of replicas for each volume and the replica reserve usage. If more than the desired number of replicas exist, consider decreasing the replica reserve percentage. If fewer than the desired number of replicas exist, consider increasing the replica reserve percentage. A low free replica reserve can indicate optimal use of replica reserve space, if the desired number of replicas exist.

Guidelines for Sizing Delegated Space

The secondary group administrator delegates space to the primary group when configuring the group as a replication partner. The secondary group administrator can modify the partner configuration and increase or decrease delegated space as needed.

Ideally, you should request from the secondary group administrator only enough delegated space to store the desired number of replicas for each volume.

NOTE: If your system has delegated space configured across multiple storage pools, the size of the space in at least one of the pools must be greater than the volume size. Otherwise, replications will fail.

For example, if you have 4 pools with 20GB of space each, but the volume size is 30GB, one or more of the pools must be changed to greater than 30GB for replications to succeed.

Use the following guidelines for sizing delegated space:

- Recommended value

  Add together maximum replica reserve space requirements for all primary group volumes you want to replicate to the secondary group and request at least that much delegated space.

  If you later decide to replicate additional volumes, the secondary group administrator might need to increase the delegated space.
- Space-efficient value

  You might want to request delegated space that is less than the recommended value.

  Initially, replica reserve is not fully allocated. Instead, it increases automatically, based on volume usage. This increase enables you to overprovision replica reserve (provision more space than what is physically available). When you overprovision replica reserve, the total maximum replica reserve space for all the partner volumes exceeds delegated space.

CAUTION: If you overprovision replica reserve, a volume's replica reserve might not be able to increase automatically or through administrative action, preventing the replication operation from completing.

For example, you are replicating five volumes, and the maximum combined replica reserve needed is 70GB. If you allocate 50GB for delegated space, replica reserve is overprovisioned by 20GB. If you specify 70GB for delegated space, each volume's replica reserve can increase to its maximum.

After you set up replication, you should monitor delegated space usage. If free delegated space is low and the replica volume reserve for each replicated volume has not reached its maximum, consider increasing the delegated space.

How Volume Changes Affect Replication Space

How much space you need for replication depends on the volume size and the rate of volume changes.

The first replication of a volume copies the entire volume contents from the primary group to the secondary group. Subsequent replication operations transfer only the data that changed since the previous replication. Replication time and space requirements increase as the amount of transferred data increases.

It can be difficult to estimate the rate of volume changes because volume usage can vary. Therefore, it can be difficult to estimate replication time and space requirements. For example:

- Although some applications perform a consistent number of volume writes, others have a workload that changes daily. Therefore, one replication operation might transfer little data and complete quickly, while another replication might transfer a large amount of data and take a long time.
- In some cases, a volume might appear to have few changes, but the transferred data is relatively large. Random writes to a volume can result in a large amount of transferred data, even if the actual data changes are small.
- Some disk operations, such as defragmenting a disk or reorganizing a database, can increase the amount of transferred data. However, the defragmentation or reorganization can make subsequent replications more efficient.

In addition, because volume usage can change over time, replication space that was adequate for one workload might become inadequate when you add more users.

If a replication operation requires copying a large amount of data, you might want to use manual transfer replication.

For each replication operation, you can display the amount of data that the primary group is transferring to a replication partner. You can also display the replication history for a volume and the amount of data transferred for each replication operation.

About Replication Partners

Before you can replicate volume and NAS container data between two PS Series groups, you must configure the groups as replication partners.

Each partner plays a role in the replication of a volume, and you can monitor replication activity and manage replicas from either partner:

- Primary group

  Location of the volume. The primary group administrator configures the secondary group as a replication partner and initiates the replication operation. Replication of the volume is considered outbound from the view of the primary group.
- Secondary group

  Location of the volume's replica set. The secondary group administrator configures the primary group as a replication partner and provides space for replicas. Replication of the volume is considered inbound from the view of the secondary group (sometimes called the destination group).

Mutual authentication using passwords provides security between partners. You can test the replication partnership before you begin replication. If you do not specifically test the partnership, it will not be verified until you begin replication.

Partners use port 3260 for replication activity. After you configure the replication partners, you can replicate a volume or replicate all the volumes in a volume collection.

Replication Partner Requirements

To be replication partners, the two groups must meet the following requirements:

- The primary group must have enough free space for the local replication reserve for each replicated volume. Local replication reserve is located in the same pool as the volume. See About Local Replication Reserve.
- The secondary group must have enough free space to delegate to the primary group. See Guidelines for Sizing Delegated Space.
- The groups must have network connectivity. The link between groups must support full IP routing and must provide sufficient bandwidth to complete replication operations in a timely manner.
- The network link between the groups must be secure (for example, through use of a firewall, VPN, or encryption).
- The groups must run compatible versions of the PS Series firmware, as detailed in the firmware matrix below. If the groups are not running the same firmware version, features in the most recent firmware version might not be available. In some cases, you must also disallow firmware downgrades. Table 48. Replication Support – Partner Firmware Matrix shows the firmware versions compatible with version 9. You can also check the *Dell EqualLogic PS Series Storage Arrays Release Notes* for the most current list.

Table 48. Replication Support – Partner Firmware Matrix

Firmware on Group	Firmware on Partner
v9.1.x	v9.0.x, v9.1x
v9.0.x	v8.0.x, v8.1.x, v9.0.x
v8.1.x	v8.0.x, v8.1.x
v8.0.x	v7.0.x, v 7.1.x, v8.0.x, v8.1.x
v7.0.x, v7.1.x	v6.0.x, v7.0.x, v7.1, v8.0.x
v6.0.x	v5.1.x, v5.2.x, v6.0.x. v7.0.x, v7.1.x
v5.1.x, v5.2.x	v5.0.x, v5.1.x, v5.2.x, v6.0.x
v5.0.x	v5.0.x, V5.1.x, v5.2x

Primary Group Replication Attributes

Table 49. Primary Volume Replication Attributes describes attributes that you set when you configure a volume for replication in the primary group. You can modify the replication configuration and change the attribute values.

Table 49. Primary Volume Replication Attributes

Attribute	Description
Replication partner	Partner that stores the volume replicas. The partner must have space delegated to the group.
Local replication reserve percentage	Space for use during replication and, optionally, for storing the failback snapshot. This space is consumed from the same pool as the volume.
Local replication borrow space setting	Enabled by default. Can be disabled only from the CLI.
Failback snapshot setting	Enables you to keep the failback snapshot in the local replication reserve. The failback snapshot can expedite failback operations.
Replica reserve percentage	Portion of delegated space on the partner for storing replicas based on the volume usage.

Replication Partner Attributes

You can modify the name, group IP address, amount of delegated space and its pool, passwords, and contact information for a replication partner.

NOTE:

- Replication partner changes you make on the secondary group are not updated on the primary group until the next replication.
- On the replication partner, if you create a container with an SMB share and use a container name that already exists, enabling read-only access on the replica container fails, but no indication of the failure is returned. An event is generated in the event log indicating that the failure was due to a duplicate share name.

When you configure a replication partner, you specify values for the attributes described in Table 50. Replication Partner Attributes. You can also modify the partner configuration and change the attribute settings as needed.

Table 50. Replication Partner Attributes

Attribute	Description
Group name and group IP address	Name and IP address of the group that you want to configure as a replication partner. The group name is case sensitive and can accept up to 54 ASCII characters.
Description	Optional description for the partner. Descriptions can be up to 63 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
Contact information	Optional contact information for the partner administrator: Name — Up to 63 characters. Email address — Up to 31 characters. Phone numbers — Up to 31 characters. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
Two passwords	Passwords for mutual authentication. Each partner supplies a password to the other partner, which validates the password: the inbound password on one partner is the outbound password on the other, and vice versa.

Attribute	Description
	Passwords are case sensitive and can include up to 254 ASCII characters.
Delegated space	Amount of space to delegate to the partner. Required only if the group stores replicas from the partner. See About Delegated Space and Replica Reserve.
NAS configuration	If you are replicating NAS containers, the replication partner must have a compatible NAS configuration.

Requirements for NAS Container Replication Partners

NAS container replication partners must meet the following requirements:

- Both inbound and outbound replication must be enabled for both replication partners.
- Source and destination NAS clusters must be running the same FS Series firmware versions.

NOTE: If you partner two clusters for replication, one with antivirus enabled and the other without antivirus enabled, any SMB shares that were configured on the original container with antivirus enabled are not restored to the promoted replica that does not have antivirus enabled. You will have to manually create the same shares that were on the original on the promoted replica container to be able access their data on the replica container.

You can replicate containers between clusters running different hardware. For example, you can replicate from an FS7500 cluster to an FS7610 cluster.

You cannot use the Manual Transfer Utility to facilitate NAS container replication.

NOTE: The requirements for NAS container replication are not validated when you create the replication partnership.

Configure a Volume for Replication

When you have configured at least one replication partner that has delegated space to the group, you can configure volumes for replication.
For information about configuring a replication partner, see Configure a Replication Partner.

The following prerequisites and considerations apply to volume replication:

- Thin clones — You cannot configure a thin clone for replication until you replicate the template volume to which the thin clone is attached.
- Template volumes — A template volume is read-only and cannot be failed back from the secondary group. Keeping the failback snapshot is not necessary for this type of volume.
- Gather the volume replication attributes.
- You cannot enable synchronous replication on a volume for which traditional replication is configured, and you cannot enable traditional replication on a volume for which synchronous replication is configured.

To configure a volume for replication:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Configure replication** to open the Configure Volume Replication wizard.
4. Provide the requested information in each step of the wizard and click **Next**.

   Refer to the online help resources that are available in the user interface to view descriptions of individual fields.
5. When you complete the final wizard step, review the information and click **Finish**.

After you finish the volume replication configuration, you can choose to create a replica (see Create a Replica). To verify that a replica has been created:

1. Click **Replication** and select a partner name.

2. Expand **Volume Replication** and then select **Inbound Replicas**.

   The **Inbound Replicas** panel displays information about the replicas.

You can also choose to perform the replication by using manual transfer replication. See the *Dell EqualLogic Manual Transfer Utility Installation and User's Guide* for more information.

Replication Partner Fields

You need the information in the following table for both primary and secondary groups.

Data	Description	Example
Name	Name of the primary group. Group names are case sensitive.	Nas22
IP Address	IP address of the primary group.	12.111.142.11
Description	(Optional) A description of the partner, such as its location.	DT Site
Administrator	(Optional) Name of the administrator for the primary group.	Mike Smith
Email	(Optional) Email address of the administrator for the primary group.	msmith@dtinfo.com
Telephone Number	(Optional) Telephone number of the administrator for the primary group.	613-555-3456
Passwords	You must specify mutual passwords.	123abc123
Pool	(Optional) Storage pool used for replication.	RepPool
Delegated Space	Amount of storage space the group delegates to the partner.	35GB

NOTE: You will not see password or configuration problems between partners until you enable replication on a volume. If you receive a login error message, make sure that you specified the correct passwords when configuring the partners.

Modify a Partner Group Name or IP Address

To modify a partner group name or IP address:

1. Click **Replication** and then select the partner name.
2. In the Activities panel, click **Modify partner settings**.
3. Change the group name, IP address, or description.
4. Click **OK**.

NOTE: Group names can be up to 54 ASCII characters. Descriptions can be up to 63 ASCII characters.

Modify Partner Passwords

If you make a modification on one partner, you must make the reciprocal modification on the other partner. The password in the **Password for partner** field on one partner must match the password in the **Password obtained from partner** field on the other partner.

To modify partner passwords:

1. Click **Replication** and then select the partner name.
2. In the Activities panel, click **Modify passwords** to open the Modify Passwords dialog box.
3. Change the passwords as needed.
   - The Inbound password is used by the partner to provide access to the current group.
   - The Outbound password is used by the current group to gain access to the replication partner.

   NOTE: The passwords fields accept up to 254 bytes of data.
4. Click **OK**.

Modify Partner Contact Information

To modify partner contact information:

1. Click **Replication** and then select the partner name.
2. Click **Modify partner settings**.
3. Change the contact name, email address, or phone numbers.
4. Click **OK**.

Manage Space Delegated to a Partner

You can modify the space delegated to a partner, subject to the following restrictions:

- You cannot decrease the space delegated to a lower capacity than is currently reserved for the partner's replicas.
- The Pool Space table shows how pool space is currently used and how much space is used after the change.

  If more than one pool is configured to have delegated space, the first row of the table shows the overall statistics for delegated space usage across all pools. The lower rows show this information for the currently selected pool.

  If the new delegated space exceeds the capacity of a pool, the color of the table cell showing free pool space changes to red.

For additional information, see Guidelines for Sizing Delegated Space.

To modify the space delegated to a partner:

1. Click **Replication** and expand the replication partner.
2. Click **Volume Replication**.
3. In the Activities panel, click **Manage delegated space** to open the Manage Delegated Space dialog box.
4. Enter the new delegated space value. (If the group has more than one pool, you can specify a new value for each one as needed.)
5. Click **OK**.

Modify Replication Services

To modify the replication services on a partner that is being replicated (the source), and enable or disable volume or NAS container replication:

1. Click **Replication** and select the partner name.
2. Click **Manage replication services** to open the Modify Replication Services dialog box.

   NOTE: This option appears in the Activities panel only if you can use the replication partner for both volume and NAS container replication.
3. Enable or disable volume or NAS replication.
4. Click **OK**.

Delete a Replication Partner

Deleting a replication partner breaks the replication relationship between the two groups. The next replication of a volume configured to use the deleted partner pauses or fails.

Deleting a partner deletes any inbound replicas stored in space that the group delegated to the partner. The delegated space becomes free pool space. Any replicas stored on the deleted partner are not deleted, and you can access them by logging in to the partner.

NOTE: If the group is hosting a recovery volume from the partner, before you delete the partner either:

- **Demote the recovery volume to an inbound replica set (which is deleted when you delete the partner). Double-click the recovery volume in the far-left panel and click Demote to replica set.**
- **Promote the recovery volume to a permanent volume.**

To delete a replication partner:

1. Determine whether NAS container replication is enabled. If so, delete all NAS container replication relationships.
2. Determine whether the group is hosting a recovery volume from the partner. If not, go to step 4.
3. For hosted recovery volumes, either:
   - Demote the recovery volume to an inbound replica set (which is deleted when you delete the partner). Double-click the recovery volume in the far-left panel and click **Demote to replica set**.
   - Promote the recovery volume to a permanent volume.
4. Click **Replication** and select the partner.
5. Click **Volume Replication**.
6. Click **Pause inbound**.
7. Select the partner.
8. Click **Delete partner**.
9. Confirm that you want to delete the partner.

Modify Volume Replication Configuration Settings

NOTE: The space currently used to store replicas represents the lower limit for replica reserve. You cannot decrease replica reserve below this limit. Changes are not applied until the next replication.

1. Click **Volumes**.
2. Expand **Volumes** and then select a volume that is configured for replication.
3. Click **Modify replication settings** in the Activities panel. Note that this link will not be visible if the selected volume is not configured for replication.
4. Click the **General** tab to open the Modify Volume Replication Settings dialog box.
5. Modify the fields as needed.

   Refer to the online help resources that are available in the user interface to view descriptions of individual fields.

   NOTE: If you select Keep failback snapshot, you must create a replica to establish the failback snapshot.

Pause and Resume Replication to or from a Partner

You can pause and then resume replication as needed. Some operations require temporarily pausing replication.

1. Click **Replication**.
2. Select the replication partner on which you want to manage replication.
3. In the **Activities panel**, select either:
   - Pause outbound — Halts current and future outbound replication to the selected partner.
   - Pause inbound — Halts current and future inbound replication from the selected partner.

   In each case, you are prompted to confirm your decision to pause.
4. In the Activities panel, select either:
   - Resume outbound — Reenables the ability to perform outbound replication.
   - Resume inbound — Reenables the ability to perform inbound replication.

Pause and Resume Replication of a Volume

You can pause and resume volume replication. For example, tasks such as promoting a replica set require you to first pause volume replication.

To pause replication for a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Pause volume replication**.
4. When prompted to confirm the decision, click **Yes**.

To resume replication for a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. In the Activities panel, click **Resume volume replication**.
4. When prompted to confirm the decision, click **Yes**.

Pause and Resume Outbound Replication

To pause outbound replication to a partner:

1. Click **Replication** and then expand the partner name.
2. Click **Volume Replication** → **Outbound Replicas** → **Pause outbound** → **Yes**.

To resume outbound replication to a partner:

1. Click **Replication** and then expand the partner name.
2. Click **Volume Replication** → **Outbound Replicas** → **Resume outbound**.

Pause and Resume Inbound Replication

You can pause and resume inbound replication when it is necessary to suspend replication temporarily for other operations.

To pause inbound replication from a partner:

1. Click **Replication**.
2. Expand the partner name.
3. Select **Inbound Replicas**.
4. In the Activities panel, click **Pause inbound**.

To resume inbound replication from a partner:

1. Click **Replication**.
2. Expand the partner name.
3. Select **Inbound Replicas**.
4. In the Activities panel, click **Resume inbound**.

Disable Replication

Disabling replication for a volume or volume collection unconfigures replication on the volume or volume collection.

Disabling replication does not delete the volume replicas stored on the partner. You can log in to the partner and access the replicas.

The following considerations apply:

- If you later reconfigure replication on the same volume to the same partner, you must delete the existing replica set on the partner before creating a replica. The first replication is a complete copy of the volume data.

- When you disable replication on a volume, the delegated space on the secondary group that is storing the replicas becomes an *unmanaged space*. You cannot manage this space from the primary group. If you do not need the replicas, log in to the secondary group and delete the replica set.
- You cannot disable replication on a template volume if any attached thin clones have replication enabled.

To disable replication for one volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Disable volume replication**.
4. When prompted, click **Yes** to confirm that you want to disable replication on the volume.

To disable replication for a volume collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection name.
3. Click **Disable collection replication**.
4. When prompted, click **Yes** to confirm that you want to disable replication on the volume collection.

Cancel Volume Replication

You can cancel an in-progress volume replication.

NOTE: To temporarily stop volume replication instead of canceling it, pause the replication.

1. Click **Volumes**.
2. Expand **Volumes** and then select a volume that is in the process of creating a replica.
3. Click **Cancel replica creation** in the Activities panel.

Replicate Volume Collections

Volume collections enable you to perform an operation on multiple volumes at the same time.

If you replicate a volume collection, the resulting set of replicas is called a replica collection. When complete, a replica collection contains one replica for each volume in the collection. A replica collection set is the set of all the replica collections for a volume collection.

The following prerequisites and considerations apply when replicating volume collections:

- Configure the individual volumes in the collection for replication to a partner.
- Configure the volume collection to replicate to the same partner.
- If the volume is large, consider using the Manual Transfer Utility (MTU) to perform the initial upload of data to the partner.
- You should always monitor replication activity to make sure that the replication operations complete as expected.

To replicate a volume collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection.
3. Click **Create replica** to open the Create Replica Collection dialog box.
4. (Optional) Select **Perform manual replication**.
5. Click **Yes**to start the replication.

   NOTE: If replication operations are taking longer than expected, make sure that you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link can cause long replication times.

Configure a Volume Collection for Replication

You can simultaneously replicate data in related volumes by replicating the volume collection. The resulting set of replicas is called a replica collection.

NOTE: To replicate a volume collection, you must configure all the volumes in the collection to replicate to the same partner.

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection that you want to replicate.
3. Click **Configure replication** to open the Configure Replication Settings dialog box.
4. Select the replication partner for the volume collection and make sure each volume is configured to replicate to the selected partner:
    - Click the **not replicated** link if a volume is not configured for replication, and configure the volume.
    - Click the partner name link if a volume is configured to replicate to a different partner. Modify the volume replication configuration and change the partner.
5. Click **OK**.

Display Replication Activity and Replicas for a Volume Collection

To display replication activity and replicas for a volume collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection name.
3. Click the **Replication** tab.

The **Replication Summary** panel shows:

- Replication partner for the collection
- Replication schedules for the volume collection, including the next scheduled replication operation, if any

The **Remote Replicas** panel shows the replica collections for the volume collection. Expand a replica collection to see the individual replicas and their status.

Modify Volume Collection Replication Configuration Settings

You can modify the replication configuration of a volume collection or of the volumes in the collection.

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection that you want to modify.
3. Click **Modify replication settings** to open the Volume Collection – Modify Replication Settings dialog box.
4. To modify the settings, double-click the volume in the collection list to open up the replication settings for that volume. Make adjustments for each volume as necessary.
5. Click **OK** to save your changes.

About Replicas

Similar to a snapshot, a replica represents the contents of a volume at the time the replica was created. Each replicated volume has a replica set, which is the set of replicas created over time. You can create replicas of individual volumes or volume collections. You can create replicas manually at any time, or you can set up a schedule.

Individual replicas are identified by the date and time that the replication operation started. The replica set name is based on the volume name and includes a dot-number extension to ensure that all replica set names are unique, in case two different partners replicate volumes with the same name to the same group. The number in the extension reflects the order that each partner was configured as a replication partner to the group. For example, all replica sets from the first configured partner have a dot-1 extension (such as `dbase.1`). Replica sets from the next configured partner have a dot-2 extension (such as `dbase.2`).

A volume and its replica set are always stored in different groups connected by a robust network link. Separating the groups geographically protects volume data in the event of a complete site disaster.

All replicas are thin provisioned by default.

Create a Replica

The first time that you replicate a volume to a partner, the primary group copies the entire volume contents to replica reserve on the secondary group. Subsequent replication operations transfer only the volume data that changed since the previous complete replication.

Before you initiate replication, be aware of the following considerations:

- Very large data transfers might exceed the capacity of the network link between the primary group and the secondary group. For replication operations that require transferring large amounts of data, consider using manual transfer replication.
- You need to set up the replication partners and configure the volume or volume collection for replication.
- Template volumes have the following constraints:
  - You must replicate a template volume before replicating any of its thin clones.
  - You can replicate a template volume only one time.
  - While a template volume replication is in progress, the inbound template volume replica set on the secondary group does not use the correct template volume icon (blue cylinder). When the replication completes, the correct icon appears.

To create a replica:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click **Create replica** to open the Create Replica dialog box.
4. (Optional) Select **Perform manual replication** to use manual transfer replication to perform the replication.
5. Click **Yes** to start the replication.

Monitor replication to make sure that replicas complete in a timely manner. To monitor replication activity:

1. Click **Volumes**.
2. Expand **Volumes** and select a volume.
3. Click the **Replication** tab to display the Replication Summary and Remote Replicas panels.

If replication operations take longer than expected, make sure that you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link can cause long replication times.

Delete Outbound Replica Sets or Replicas

You can delete unwanted outbound replica sets or replicas.

NOTE: Deleting a replica set disables replication on the volume. If you reenable replication on the volume, the first replication is a complete transfer of volume data. To delete the most recent replica, you must disable replication for the volume.

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume name.
3. Click the **Replication** tab.
4. In the Remote Replicas panel, select the **Volume replicas** view.
5. Delete outbound replicas or replica sets as follows:
   - To delete an entire replica set, select the replica set and click **Delete replica set** in the Activities panel, or right-click the replica set and select **Delete Replica Set** from the menu.
   - To delete a single replica, select the replica, and click **Delete replica** in the Activities panel, or right-click the replica and select **Delete Replica** from the menu.

- To selectively delete replicas from a set, select the replica set and click **Delete replicas** in the Activities panel, or right-click the replica set and select **Delete Replicas** from the menu. In the dialog box that opens, you can select the replicas to be deleted. Hold down the `Control` key while clicking to select multiple replicas, then click **OK** to delete the selected replicas.

Delete Inbound Replica Sets or Replicas

If the primary group is not available, you can delete replicas and replica sets when you are logged in to the secondary group. However, if you delete replicas or replica sets from the secondary group, the primary group information is not updated and errors can result.

The following considerations apply:

- Dell recommends that you delete replicas when you are logged in to the primary group.
- Deleting a replica set disables replication on the volume.
- To delete the most recent replica, you must disable replication for the volume.

To delete an inbound replica or replica set:

1. Click **Replication**.
2. Expand the partner name.
3. Expand **Volume Replication**.
4. Click **Inbound Replicas**.
5. You must pause inbound replication before you delete an inbound replica set. In the Activities panel, click **Pause inbound**. When prompted, click **Yes** to confirm the operation.
6. Delete inbound replicas or replica sets as follows:
   - To delete an entire replica set, select the replica set and click **Delete replica set**, or right-click the replica set and select **Delete replica set** from the menu.
   - To delete a single replica, select the replica and click **Delete replica**, or right-click the replica and select **Delete replica** from the menu.
   - To selectively delete multiple replicas from a set, select the replica set and click **Delete replicas**, or right-click the replica set and select **Delete replicas** from the menu. In the Delete Replicas dialog box, hold down the `Control` key while selecting the replicas to delete, then click **OK**.
7. In the Activities panel, click **Resume inbound** to resume inbound replication.

The replica or replica set no longer appears in the group. However, the replica or replica set still appears on the partner (primary group), if it is available. You can log in to the primary group and delete the replica or replica set.

Delete Outbound Replica Collection Sets, Replica Collections, or Replicas

You can delete an outbound replica collection, an outbound replica that is part of a replica collection, or the entire replica collection set for a volume collection. Deleting a replica collection deletes all the replicas that are in the collection.

The following considerations apply:

- If you delete a replica from a replica collection, the replica collection no longer represents the contents of the volumes in the collection at the time you created the replica collection.
- Deleting a replica collection set disables replication on the volume collection. If you reenable replication on the volume collection, the first replication of each volume is a complete transfer of volume data.
- To delete the most recent replica, you must disable replication for the volume.

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection name.
3. Click the **Replication** tab.
4. Delete items as follows:
   - To delete multiple collections or replicas, select the replica set and click **Delete Replica Collections**. In the Delete Replica Collections dialog box, select the collections or replicas to delete and click **OK**. (Hold down the `Control` key to select multiple items.)

- To delete a replica collection set, select it in the Remote Replicas panel, then click **Delete replica collection set**.
- To delete a replica collection, expand the replica collection set, then select the replica collection and click **Delete replica collection**.
- To delete a single replica from a replica collection, expand the replica collection, then select the replica and click **Delete replica**.

Delete Inbound Replica Collection Sets, Replica Collections, or Replicas

You can delete an inbound volume replica collection, an inbound replica that is part of a replica collection, or the entire inbound replica collection set for a volume collection. Deleting a replica collection deletes all the replicas that are in the collection.

The following considerations apply:

- Dell recommends that you delete replicas when you are logged in to the primary group.
- Deleting a replica collection set disables replication on the volume collection.
- To delete the most recent replica, you must disable replication for the volume or collection.

1. Click **Replication** and expand the partner name.
2. Expand **Volume Replication**.
3. Click **Inbound Replica Collections**.
4. Click **Pause inbound** if you are deleting an inbound replica collection set. When prompted, click **Yes** to confirm the operation.
5. Delete items as follows:
   - To delete a replica collection set, select the replica set and click **Delete replica collection set**.
   - To delete multiple replica collections, select a replica collection set and click **Delete Replica Collections**. When the dialog box opens, hold down the `Control` key while selecting replica collections and then click **OK**.
   - To delete a single replica collection, expand the replica collection set, select the replica collection, and click **Delete replica collection**.
   - To delete an individual replica, expand the replica collection, select the replica, and click **Delete replica**.
6. In the Activities panel, click **Resume inbound** to resume inbound replication.

The replica, replica collection, or replica collection set no longer appears in the group. However, the replica or replica set still appears on the partner (primary group), if it is available. You can log in to the primary group and delete the replica or replica set.

About Replication Borrowing

Replication borrowing allows you to temporarily increase the available replication space for a volume by borrowing space from other sources. By allowing replicas to borrow beyond their configured reserve, this available space is utilized more efficiently and the max-keep policy is maintained.

Keep the following considerations about replication borrowing in mind:

- All members in a group must be running PS Series firmware version 8.0 or later. Replication borrowing does not work in environments running mixed versions of the firmware.
- Replication borrowing is always enabled and happens automatically.

Example of Replication Borrowing

Replication requires a replica reserve that is sufficiently large enough to hold all of the replicas that you want to keep. For a volume that is scheduled to replicate once per day with a max-keep of 3, if the replica reserve is set too high, a lot of space will not be used. This underutilization of space might lead you to reduce the reserve to hold only three days of replication data.

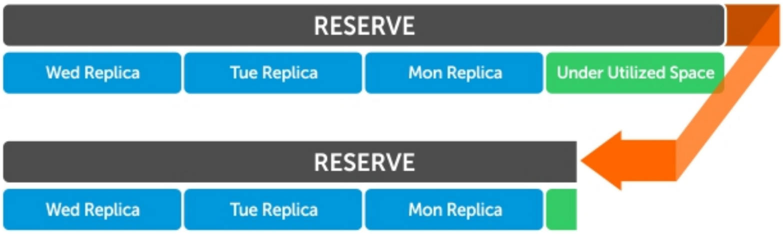

Previous to v8.0, writing a significant amount of data to a volume might result in the day's replica being unusually large. The size of this replica might be so large that older replicas are deleted from the replica set, which means that the replica set no longer retains three replicas.

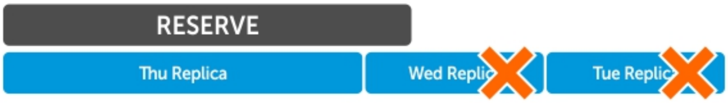

As of v8.0, replication borrowing allows replica sets to borrow enough space to hold the replicas that would otherwise be deleted. As long as the pool has enough unused space to borrow, the replica set can borrow space to maintain the max-keep policy.

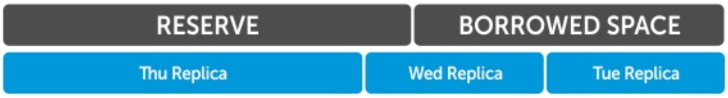

After a few days of normal I/O activity, the max-keep policy will eventually delete the large replica for which space needed to be borrowed. The remaining replicas will probably fit within the replica reserve, and the replica set might not need to borrow space any longer.

About Cross-Platform Replication

Cross-platform replication provides a method for performing asynchronous volume replication between the PS Series groups and Dell Storage Centers. This feature preserves the functional and operational models embodied in the current replication implementation for each storage system. Dell Storage Manager (DSM) is the management tool for Storage Centers and it also supports managing the PS series. DSM should be used to manage cross-platform replication as Group Manager supports read-only view of cross-platform replication.

See the *Dell Storage Manager Administrator's Guide* for detailed information about cross-platform replication using DSM.

Unsupported Functionality

Cross-platform replication does not support the following operations and components:

- Replication management operations using Group Manager. Use DSM for these operations.
- Manual Transfer Utility (MTU)
- Host Integration Tools
- Replication of VMware virtual volumes (VVols)
- Replication of volume collections and NAS containers. Cross-platform replication supports volume replication only
- Replication of thin clones and template volumes
- Shrinking of a volume that is configured for replication if the source partner is an EqualLogic group, and the destination partner is a Compellent array

Limitations

- Group Manager no longer identifies replica sets on the destination partner that do not have a corresponding volume on the source partner. DSM provides a direct view of both partners to the same administrator.
- When using cross-platform replication, volumes are automatically thin-provisioned when you promote the replica set.
- Source Volume's (network address authority) NAA IDs are not preserved on the replication volume when using cross-platform replication.
- Downgrading to an earlier version of the firmware is not allowed if a cross-platform replication partnership has been configured.
- PS group volumes with a sector size of 4K bytes cannot be replicated to a Storage Center.

About Schedules

You can create schedules to automatically perform volume and NAS container operations at a specific time or on a regular basis (for example, hourly or daily). For example, you can create a schedule to create snapshots or replicas of a volume, volume collection, or NAS container.

The following restrictions apply:

- If a volume is part of a volume collection, make sure that a schedule for the collection does not overlap a schedule for the volume.
- You cannot use manual transfer replication with a replication schedule. Schedules apply only to network replications. Scheduled replications do not run until any in-process manual transfer replications are complete.
- If you are creating a replication schedule, you must configure both groups to be replication partners for each other.
- You cannot schedule a snapshot or replication operation for a template volume.

Using a schedule can generate many snapshots or replicas, so make sure that you have sufficient snapshot or replication space. You can set a limit on the number of snapshots or replicas that a schedule creates.

Create a Schedule

1. Click **Volumes**, expand **Volumes**, and then select the volume name.
2. In the Activities panel, click **Create schedule** to start the wizard.
3. In the **Create schedule** wizard:

    a. Type the schedule name. The name can be up to 63 bytes and is case-insensitive. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

    b. Turn on the schedule by selecting the **Enable schedule** checkbox. You can defer this action if required.

    c. Select the type of schedule.

    d. Select the frequency for the schedule, or select **Reuse existing schedule** to use an existing schedule as a model for the new schedule

    e. Click **Next**.

    Depending on the schedule option that you selected in step 2d, one of the following dialog boxes opens:

    - **Run once** opens Create Schedule – Run Once
    - **Hourly schedule** opens Create Schedule – Hourly Schedule
    - **Daily schedule** opens Create Schedule – Daily Schedule
    - **Reuse existing schedule** opens Create Schedule – Select Existing Schedule
4. In the dialog box, select or type the following settings:

    - When and how often the schedule runs
    - Number of snapshots or replicas to keep
    - Whether snapshots are read-write or read-only

    If you selected the **Reuse existing schedule** option, you will first need to select a schedule from the list.
5. Click **Next** and review the information in the summary to determine whether the schedule configuration is correct.
6. Click **Finish**, or click **Back** to make changes.

Modify a Schedule

NOTE: NAS containers require a different scheduling procedure. See Modify a NAS Container Snapshot Schedule.

You can modify any attribute except for:

- Schedule type (snapshot or replication)
- Schedule frequency (one time, hourly, or daily)

Attributes that are available for modification depend on the type and frequency of the schedule.

To modify a schedule:

1. Click **Volumes**, and then either:
   - Expand **Volumes**, select the volume, and click the **Schedules** tab.
   - Expand **Volume Collections**, select the collection, and click the **Schedules** tab.
2. Select the schedule in the Snapshot and Replication Schedules panel and click **Modify**.
3. Change the schedule attributes as needed:
   - Schedule name – Name can be up to 63 bytes and is case insensitive. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - Start and end dates
   - Time of day and repetition
   - Snapshot settings (keep, writable)
   - Enable or disable the schedule
4. Click **OK**.

Delete a Schedule

You can delete unwanted schedules for volumes or collections. Deleting a schedule does not affect the existing snapshots or replicas that the schedule created.

To delete a schedule:

1. Click **Volumes**, then either:
   - Expand **Volumes**, select the volume, and click the **Schedules** tab.
   - Expand **Collections**, select the collection, and click the **Schedules** tab.
2. Select the schedule in the **Snapshot and Replication Schedules** panel.
3. Click **Delete**.
4. Confirm that you want to delete the schedule.

Create a NAS Container Snapshot Schedule

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. In the Activities panel, click **Create schedule** to open the wizard.
3. Provide the requested information in each step of the wizard and click **Next**.
4. Either:
   - Click the **Copy** link to copy the schedule information so that you can paste it into a text file and save it for future reference.
   - Click **Back** to make changes.
   - Click **Finish** if the schedule is correct.

Modify a NAS Container Snapshot Schedule

1. Click **NAS**, expand **NAS Cluster** and **Local Containers**, and then select the NAS container name.
2. Click the **Schedules** tab.
3. In the Snapshot Schedules panel, select the schedule name and click **Modify schedule**.
4. In the **Modify schedule** dialog box, modify the configuration settings for the schedule as needed.

NOTE: You cannot change the schedule type.

5. Click **OK**.

Monitor NAS Snapshot Schedules

1. Select **Monitoring** in the navigation menu.
2. Under NAS Schedules, select **NAS Snapshot schedules**.
3. Right-click any of the data fields to modify, delete, enable, or disable a schedule:
   - **Name** — User-defined schedule name.
   - **Container** — Name of the source NAS container. Click the name to go to the container's data.
   - **Create** — Schedule action (in this case, create a snapshot).
   - **Run** — Date on which the schedule started.
   - **Time** — Time of day that the schedule runs.
   - **Status** — Whether the schedule is enabled (running) or disabled (paused).
   - **Actions** — Tasks that you can perform for this schedule (including modify, enable or disable, and delete).

Disable a NAS Container Snapshot Schedule

To temporarily stop a schedule from creating NAS container snapshots, you can disable the schedule. Enable the schedule to resume the regular creation of snapshots.

To disable a NAS container snapshot schedule:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Schedules** tab.
3. In the Snapshot and Replication Schedules panel, right-click the schedule that you want to disable and click **Disable schedule**.
4. Click **Yes** to confirm that you want to disable the schedule.

Delete a NAS Container Snapshot Schedule

Deleting a snapshot schedule for a NAS container does not affect any of the snapshots created by the schedule.

To delete a snapshot schedule for a NAS container:

1. Click **NAS**, expand **NAS Clusters** and **Local Containers**, and then select the NAS container name.
2. Click the **Schedules** tab.
3. In the Snapshot and Replication Schedules panel, select the schedule and click **Delete**.
4. Click to confirm that you want to delete the schedule.

17

About Data Recovery

A PS Series group is part of a comprehensive backup and data protection solution.

Snapshots provide quick recovery and offloading backup operations. Restore operations are more reliable because snapshots ensure the integrity of the backed-up data.

Replication protects data from serious failures such as destruction of a volume during a power outage, or a complete site disaster.

The following list provides more information:

- Snapshots – You can use a snapshot to restore a volume to the state it was in at the time the snapshot was created. In addition, you can make snapshots available to iSCSI initiators or clone the snapshot to create a new volume.
- Setting snapshots online or offline – By default, a snapshot is offline. You can set a snapshot online, making it accessible to iSCSI initiators that match one of the snapshot's access control policies.If you set a snapshot offline, any current iSCSI connections to the snapshot are lost.
- Cloning snapshots to create volumes – Cloning a snapshot creates a new standard volume, template volume, or thin clone volume with a new name and new iSCSI target name, but with the same reported size, pool, and contents as the original volume at the time that you created the snapshot.

  The group allocates space equal to the volume reserve that you specify for the new volume. If you reserve snapshot space for the new volume, the group allocates additional space.

  The snapshot still exists after the clone operation.
- Restoring a volume from a snapshot – You can restore a volume from a snapshot, and replace the data in the current volume with the volume data at the time you created the snapshot. The snapshot still exists after the restore operation.
- Replicas – You can use replicas to recover data on a volume in several ways:
  - Failing over a volume to its replica, and later failing back to the primary group.
  - Promoting a replica to a recovery volume and replicating its data back to the primary group. You can later demote the volume and resume the original replication, or you can make the promotion permanent.
  - Moving a replica set to another storage pool.
  - Switching partner roles so that the primary group is now the secondary group and vice versa.
- Promoting an inbound replica set to a recovery volume – To temporarily fail over a volume (or template or thin clone) to the secondary group, you promote the inbound replica set to a recovery volume (or recovery template or recovery thin clone) and snapshots. Users can connect to the recovery volume and resume accessing the volume data.

  A recovery volume name is generated automatically, based on the volume name, with a dot-number extension (for example, vol01.1). You can choose to keep the same iSCSI target name as the original volume to facilitate iSCSI initiator connections to the recovery volume.

  Promoting an inbound replica set does not require any additional space on the secondary group, because it reduces delegated space by the size of the volume's replica reserve. Thin provisioning is automatically enabled for recovery volumes.
- Replicating a recovery volume to the primary group – When the original volume on the primary group becomes available, you can replicate the recovery volume to the primary group. This action synchronizes the data across both groups and protects the recovery volume. During the replication, initiators can continue to access the recovery volume.

  Recommendation: Dell recommends that you replicate the recovery volume to the primary group immediately before failing back to the primary group, because the volume is offline during the final replication that is part of failing back to the primary group.

  Restriction: You cannot replicate a recovery template volume.

  The Replicate to Partner operation is available only if the primary group and the secondary group are running PS Series firmware version 5.0 or greater. If you do not meet this requirement, you must perform the steps individually.

How quickly you can replicate the recovery volume depends on the presence of the failback snapshot on the primary group. The failback snapshot establishes the failback baseline, which is the point in time at which the volume on the primary group and the most recent complete replica on the secondary group have the same data. If the failback snapshot exists, only the changes made to the recovery volume are replicated. If the failback snapshot does not exist, the first replication is a complete copy of the recovery volume data.

- Moving a failback replica set to a different storage pool – You can move a failback replica set to a different pool in the primary group. Later, if you promote the failback replica set to a volume, the volume belongs to the new pool.
- Failing back to the primary group – To return to the original volume replication configuration, you can use the Failback to Primary operation.

  Recommendation: Dell recommends that you use the Replicate to Partner operation before failing back to the primary group. Although the Failback to Primary operation performs a final replication, the recovery volume is offline during the final replication.

  As part of the failback operation, a replica is created immediately on the secondary group to reestablish the failback snapshot (and set the failback baseline). Because the volume data is already synchronized between the groups, no data is actually transferred.

  Restriction: You cannot fail back a template volume. Also, the Failback to Primary operation is available only if the primary group and the secondary group are running PS Series firmware version 5.0 or greater.
- Making a temporary volume available on the secondary group – You can make a temporary copy of a volume available on the secondary group, while providing continuous access to the original volume on the primary group. Using a temporary copy is helpful when you want to perform an operation (such as a backup) on the copy with no disruption to users. When the operation completes, you can resume replicating the volume.

  NOTE: This procedure assumes that the volume does not change while available on the secondary group, or — if the volume changes — those changes are not replicated to the primary group. If you want to replicate changes, follow the procedure described in About Failing Over and Failing Back a Volume.
- Switching partner roles – You can switch the partner roles in a volume replication configuration. The original secondary group becomes the new primary group, and the original primary group becomes the new secondary group.

  NOTE: Because you cannot permanently demote a template volume, when you switch roles for a replication configuration that includes a template volume with thin clone volumes, only the thin clone replication configuration switches. Therefore, the original template volume must still exist on the original primary group after the switch, because the new thin clone replica sets depend on the template volume.
- Making promotions permanent – After promoting an inbound replica set to a recovery volume, you can make the promotion permanent, resulting in a new standard volume, template volume, or thin clone volume. You might need to perform this task if the original volume is destroyed or if you are switching roles in a replication configuration.

  NOTE: After making an inbound replica set promotion permanent, you can no longer demote the volume to the original inbound replica set.

  Restriction: Before you can make a template replica set promotion permanent, you must permanently promote all the attached thin clone replica sets.

About Recovering Data from a Snapshot

You can restore a volume from a snapshot, and replace the data in the current volume with the volume data at the time you created the snapshot. The restored volume has the same name and iSCSI target name as the original volume.

You can also clone a snapshot to create a new volume, allowing you to access the volume's data as it existed at a certain point in time without having to replace the volume's current data. Cloning a snapshot is especially useful if you need to recover individual files and roll them back to a previous revision.

In both cases, the snapshot still exists after the restore operation.

Requirements and Restrictions

- To restore a volume from a snapshot, all group members containing data from the volume or its snapshots must be online.

- You cannot use snapshots to restore data to template volumes.
- Restoring the volume from a snapshot requires taking the volume offline and terminating any iSCSI connections to the volume.
- You cannot restore a synchronous replication (SyncRep) volume from a snapshot if the snapshot's size is different from that of the volume.
- You cannot restore template volumes from snapshots.

Failback to Primary Operation (Manual)

The Failback to Primary operation consolidates multiple tasks. You can perform each task in the operation individually.

One step in the procedure requires that you create a replica. You can use manual replication if you have a large amount of data to transfer.

To fail back a volume to a partner using individual tasks:

1. From the primary group, perform the Replicate to Partner operation.
2. Log in to the secondary group to perform the remaining steps.
3. Disable any replication or snapshot schedules for the recovery volume.
4. Set the recovery volume offline.
5. Create a final replica.
6. Demote the recovery volume to the original inbound replica set:
   a. Click **Volumes**.
   b. Expand **Volumes** and then select the volume name.
   c. Click **Demote to replica set**.
7. On the primary group, promote the failback replica set to the original volume:
   a. Click **Replication**.
   b. Expand the partner name.
   c. Expand **Inbound Replicas** and select the failback replica set.
   d. Click **Promote to volume**.

Move a Failback Replica Set to a Different Pool

You can move a failback replica set to a different pool in the primary group. If you later promote the failback replica set to a volume, the volume belongs to the new pool.

To move a failback replica set to a different pool:

1. Log in to the primary group.
2. Click **Replication**.
3. Expand the replication partner and then expand **Volume Replication** (if displayed).
4. Expand **Inbound Replicas** and then select the failback replica set.
5. Click **Move replica set**.
6. Select the new pool and click **OK**.

Replicate to Partner Operation (Manual)

The Replicate to Partner operation consolidates multiple tasks. You can perform each task in the operation individually.

The following considerations apply:

- You must promote the inbound replica set to a recovery volume before performing the individual Replicate to Partner tasks.

- You can use manual replication if a large amount of data must be transferred.

  See the *Dell EqualLogic Manual Transfer Utility Installation and User's Guide* or the online help for information.

To replicate a recovery volume to a partner using individual tasks:

1. Log in to the primary group and then:

   a. Set the original volume offline.
   b. Cancel any in-progress replication.
   c. Set any snapshots for the volume offline.
   d. Demote the volume to a failback replica set: Click **Volumes**, expand **Volumes** and select the volume name, and then click **Demote to replica set**.
2. Log in to the secondary group and then:

   a. Configure the recovery volume to replicate to the primary group.
   b. Create a replica.

Switch Partner Roles Permanently

You can switch the partner roles in a volume replication configuration. The original secondary group becomes the new primary group, and the original primary group becomes the new secondary group.

NOTE: Because you cannot permanently demote a template volume, when you switch roles for a replication configuration that includes a template volume with thin clone volumes, only the thin clone replication configuration switches. Therefore, the original template volume must still exist on the original primary group after the switch, because the new thin clone replica sets depend on the template volume.

To permanently switch partner roles:

1. On the primary group:

   - Make sure that the volume replication configuration includes keeping the failback snapshot.
   - Set the volume offline.
   - Perform a final replication. This operation synchronizes volume data across the primary group and the secondary group.
2. On the secondary group:

   - Promote the replica set to a recovery volume. Make sure that you keep the ability to demote the recovery volume, in case you decide to cancel the role switch.
   - Users can now access volume data by connecting to the recovery volume.
   - Replicate the recovery volume to the primary group.
   - Make the inbound replica set promotion permanent (see Make an Inbound Replica Set Promotion Permanent).
3. On the primary group, convert the failback replica set to an inbound replica set (see Convert a Failback Replica Set to an Inbound Replica Set).

The partner role switch is complete.

Make an Inbound Replica Set Promotion Permanent

After promoting an inbound replica set to a recovery volume, you can make the promotion permanent, resulting in a new standard volume, template volume, or thin clone volume. You might need to perform this task if the original volume is destroyed or if you are switching roles in a replication configuration.

The following constraints and considerations apply:

- After making an inbound replica set promotion permanent, you can no longer demote the volume to the original inbound replica set.

- Before you can make a template replica set promotion permanent, you must permanently promote all the attached thin clone replica sets.
- You must specify:
  - A new volume name, must be unique name in the group

    Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
  - Access control credentials for the recovery volume

    Specify a CHAP user name, IP address, or iSCSI initiator name.
  - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume

To make an inbound replica set promotion permanent:

1. Click **Volumes**, expand **Volumes**, and select the recovery volume.
2. Click **Make promote permanent**.
3. Change the information in the **Convert recovery volume – Volume settings** dialog box:
   - Type a new volume name.
   - (Optional) Type a description.

   Click **Next**.
4. Specify the following information in the **Convert recovery volume – iSCSI access** dialog box:
   - CHAP user name, IP address, or iSCSI initiator name
   - Permission (either read-only or read-write)
   - Enable multihost access to targets

   Click **Next**.
5. Review the information in the **Convert recovery volume – Summary** dialog box.
6. Click **Finish**, or click **Back** to make changes.

When the operation completes, the recovery volume is converted to a volume.

Convert a Failback Replica Set to an Inbound Replica Set

The final step in the procedure for switching roles in a volume replication configuration is to permanently convert the volume's failback replica set to an inbound replica set. The following restrictions apply:

- You must perform this operation on the primary group.
- After you convert a failback replica set to an inbound replica set, you cannot promote the inbound replica set to the original volume.
- When the conversion completes, the replica set continues to appear in the Replication Partner – Inbound window, but it is no longer a failback replica set.

To convert a failback replica set to an inbound replica set:

1. Click **Replication**.
2. Expand the partner.
3. Expand **Volume Replication**.
4. Expand **Inbound Replicas**.
5. Select the failback replica set.
6. Click **Convert to replica set**.
7. Confirm that you want to convert the replica set.

Make a Temporary Volume Available on the Secondary Group

You can make a temporary copy of a volume available on the secondary group, while providing continuous access to the original volume on the primary group. Using a temporary copy is helpful when you want to perform an operation (such as a backup) on the copy with no disruption to users. When the operation is completed, you can resume replicating the volume.

NOTE: This procedure assumes that the volume does not change while available on the secondary group, or — if the volume changes — those changes are not replicated to the primary group. If you want to replicate changes, perform a volume failover and failback procedure.

To make a temporary copy of a volume available on the secondary group and then resume replicating the volume:

1. Promote the replica set to a recovery volume. Make sure that you select the option that enables you to demote the recovery volume.
2. Perform the desired operation on the recovery volume. (Be aware of recovery-volume restrictions when promoting an inbound replica set to a recovery volume.)
3. Demote the recovery volume to an inbound replica set as follows:
   a. Click **Volumes**.
   b. Expand **Volumes** and then select the recovery volume.
   c. Click **Demote to replica set**.

At this point, replication can continue as usual.

Replicate a Recovery Volume to the Primary Group

If the original volume on the primary group becomes available, you can replicate the recovery volume to the primary group. This action synchronizes the data across both groups and protects the recovery volume. During the replication, initiators can continue to access the recovery volume.

The time required to replicate the recovery volume depends on the presence of the failback snapshot on the primary group. The failback snapshot establishes the failback baseline, which is the point in time at which the volume on the primary group and the most recent complete replica on the secondary group have the same data. If the failback snapshot exists, only the changes made to the recovery volume are replicated. If the failback snapshot does not exist, the first replication is a complete copy of the recovery volume data.

The following restrictions and considerations apply:

- You cannot replicate a recovery template volume.
- Dell recommends that you replicate the recovery volume to the primary group immediately before failing back to the primary group, because the volume is offline during the final replication that is part of failing back to the primary group.
- During the procedure, you can choose to perform a required replication by using manual transfer replication.

  If you choose to use manual transfer replication, the status of the Create Replica task is in-progress until you complete the manual transfer replication. When the manual transfer replication is complete, the Replicate to Partner operation continues automatically.
- The Replicate to Partner operation is available only if the primary group and the secondary group are running PS Series firmware version 5.0 or greater. If you do not meet this requirement, you must manually perform the Replicate to Partner procedure.
- You must stop all replications of the recovery volume to the primary group before permanently promoting a replica set.

To replicate a recovery volume to the primary group:

1. Obtain the name and password for a group administrator account on the primary group.
2. Click **Volumes**.
3. Expand **Volumes** and then select the recovery volume.

4. Click **Replicate to partner** to open the Replicate Recovery Volume dialog box.
5. Specify the group administrator account name and password.
6. Select whether to perform the replication by using manual transfer replication.
7. Select whether to save the primary group administrator account name and password for future use in the current GUI session.
8. Click **OK**.

Monitor the Replicate to Partner operation to make sure that all tasks complete:

1. Open the Alarms panel at the bottom of the GUI window.
2. Click the **Failback Operations** tab. If an individual task fails, you must correct the problem and then retry the task.

When the volume demote task on the primary group completes, the original volume disappears from the list of volumes, and the failback replica set appears under Inbound Replicas in the far-left panel.

Where to Go from Here

- To create more replicas, select the recovery volume and click **Create replica**. You can also configure replication schedules on the recovery volume.
- When you are ready, you can fail back to the primary group.

Promote an Inbound Replica Set to a Recovery Volume

To temporarily fail over a volume (or template or thin clone) to the secondary group, you promote the inbound replica set to a recovery volume (or recovery template or recovery thin clone) and snapshots. Users can connect to the recovery volume and resume accessing the volume data.

A recovery volume name is generated automatically, based on the volume name, with a dot-number extension (for example, `vol01.1`). You can choose to keep the same iSCSI target name as the original volume to facilitate iSCSI initiator connections to the recovery volume.

Promoting an inbound replica set does not require any additional space on the secondary group, because it reduces delegated space by the size of the volume’s replica reserve. Recovery volumes are thin provisioned.

With some exceptions, all volume operations apply to a recovery volume.

The following prerequisites and considerations apply:

- You cannot convert a recovery template to a standard volume. You must first make the promotion permanent.
- Make sure snapshots are offline before you perform this task.
- You cannot detach a recovery thin clone.
- Some exceptions that apply to volume operations also apply to a recovery volume.
- When prompted to keep the ability to demote a replica set, make sure that you keep this ability, unless you are permanently promoting the replica set.

To promote an inbound replica set:

1. Click **Replication**.
2. Expand the replication partner.
3. Expand **Inbound Replicas** and then select the replica set.
4. Click **Promote to volume**.
5. Confirm that you want to pause inbound replication from the partner. Replication resumes automatically for all other volumes after the replica set is promoted.
6. In the **Promote replica set – Volume options** dialog box:
    - Set the volume online if you want initiators to connect to it.

- Retain the iSCSI target name of the original volume.
- Keep the ability to demote to the replica set. (Unless you are permanently promoting the replica set, make sure that you keep this ability.)

7. Click **Next** to open the Promote Replica Set – iSCSI Access panel.
8. Specify the following information:

   - Conditions that a computer must match to connect to the recovery volume. Type a CHAP user name, IP address, or iSCSI initiator name.
   - Recovery volume permission (either read-only or read-write).
   - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the recovery volume.
9. Click **Next** to open the Promote Replica Set – Summary dialog box.
10. Review the information. Click **Finish**, or click **Back** to make changes.

After the promote operation completes, the replica set disappears from the list of inbound replica sets, and the recovery volume appears in the list of volumes.

Recovery Volume Options

After you promote an inbound replica set to a recovery volume:

- You can connect to the recovery volume.
- When the original volume on the primary group becomes available, go to the next step in the failover and failback process.
- To reverse the inbound replica set promotion and cancel the failover, demote the recovery volume to an inbound replica set:

  a. Click **Volumes**.
  b. Expand **Volumes** and then select the recovery volume name.
  c. Click **Demote to replica set**.
- If the original volume becomes permanently unavailable, you can make the inbound replica set promotion permanent.

Recovery Volume Restrictions

To temporarily fail over a volume to the secondary group, you promote the volume’s inbound replica set to a recovery volume. Users can connect to the recovery volume and resume accessing the volume data.

All volume operations apply to a recovery volume, with some exceptions:

- You cannot change:

  – Volume size
  – Volume name
  – Public alias
  – RAID preference
  – Replication partner
  – Thin-provisioning settings (applicable only to recovery template volumes and recovery thin clone volumes)
  – Permission (applicable only to recovery template volumes)
- You cannot delete a recovery template volume if any of the following objects are attached:

  – Recovery thin clone volumes
  – Thin clone replica sets
  – Permanently promoted thin clone replica sets

How to Handle a Failed Operation

To check the status of a Replicate to Partner operation and the Failback to Primary operation:

1. Open the Alarms panel and click the **Failback Operations** tab.

2. Expand the recovery volume to display the status of each task in the operation.

If an individual task fails during a Replicate to Partner or Failback to Primary operation, correct the problem.

After correcting the problem, in the Failback Operations panel, right-click the failed operation and click **Retry task**. The operation continues automatically.

Fail Back to the Primary Group

When you want to return to the original volume replication configuration, you can use the Failback to Primary operation.

The following considerations apply:

- Dell recommends that you use the Replicate to Partner operation before failing back to the primary group. Although the Failback to Primary operation performs a final replication, the recovery volume is offline during the final replication.
- You cannot fail back a template volume.
- The Failback to Primary operation is available only if the primary group and the secondary group are running PS Series firmware version 5.0 or later. If you do not meet this requirement, you must perform the steps individually.
- If you choose to use manual transfer replication, the status of the create replica task remains `in progress` until you complete the manual transfer replication. When the manual transfer replication is complete, the Failback to Primary operation continues automatically.
- You must stop pending failback operations on the volume before permanently promoting a replica set.

To fail back to the primary group:

1. Obtain the name and password for a group administrator account on the primary group.
2. Click **Volumes**.
3. Expand **Volumes** and then select the recovery volume.
4. Click **Failback to primary**.
5. Confirm that you want to set the recovery volume offline to open the **Failback recovery volume** dialog box.
6. Specify the group administrator account name and password.
7. Select whether to perform the replication by using manual transfer replication.
8. Select whether to save the primary group administrator account name and password for future use in the current GUI session.
9. Click **OK**.

The failback operation immediately creates a replica on the secondary group to reestablish the failback snapshot (and set the failback baseline). Because the volume data is already synchronized between the groups, no data is actually transferred.

Volume Failover and Failback

Volume failover and failback are part of the replication recovery process.

The following constraints apply:

- You cannot replicate a recovery template volume.
- You cannot demote a template volume to a failback replica set.
- You must stop pending failback operations on the volume before permanently promoting a replica set.
- Snapshots must be offline for failback to occur.

To fail over and fail back a volume:

1. Promote the replica set to a recovery volume (and snapshots) on the secondary group, and allow initiators to connect to the volume. You can choose to keep the same iSCSI target name to facilitate iSCSI initiator access to the recovery volume.
2. When the original volume on the primary group becomes available, synchronize the volume data on both groups. Use the Replicate to Partner operation to:

    a. Demote the original volume to a failback replica set on the primary group.
    b. Replicate the recovery volume to the primary group. If you kept the failback snapshot for the original volume, only the changes made to the recovery volume are replicated.
3. When you are ready to fail back to the primary group, use the Failback to Partner operation to:
    a. Set the recovery volume offline.
    b. Perform a final replication to synchronize the volume data across both groups.
    c. Demote the recovery volume to an inbound replica set.
    d. Promote the failback replica set to a volume and snapshots. The volume represents the data that was in the most recent complete replica. The snapshots correspond to any additional replicas.

Recover Data from a Replica

If you replicate a volume to a partner, you can recover volume data from the partner. In addition, you might be able to fail over to the partner and later fail back to the original group.

If a volume is destroyed, you can fail over to the recovery group and recover data from a replica. Users can then resume access to the recovery volume. When the original volume becomes available, you can fail back to the original group.

About Data Recovery from a Replica

Effective data recovery requires a well-planned disaster-protection strategy and the regular creation of replicas and backups. To protect volume data from unrecoverable failure, you can replicate a volume to a group configured as a replication partner.

If the volume becomes unavailable, either temporarily or permanently, you can recover the data from the partner. The method of recovering data depends on the state of the groups and your specific data-recovery requirements.

When volume failure occurs, or if the primary group is unavailable because of maintenance, it is important to resume data availability as soon as possible to prevent or limit application downtime.

For example, you can clone a replica to create a new volume on the secondary group. The new volume contains the same data that existed at the time you created the replica; initiators can connect to it in the usual way. Cloning a replica has no impact on the original volume and the replication configuration. If the original volume is still available, replication can continue as usual.

In most situations where you need to recover data, the primary group is not available because of maintenance or a failure. In this case, you can temporarily or permanently fail over the volume to the secondary group and make the volume data available to initiators. If the original volume on the primary group becomes available again, you can fail back the volume to the primary group, returning to the original replication configuration.

You implement failover and failback by using the following operations:

- Promote

  Enables you to convert a replica set into a volume and snapshots. The volume contains the data represented by the most recent, complete replica. The snapshots correspond to the remaining replicas.

  For example, you can promote an inbound replica set to a recovery volume as part of a failover operation.
- Demote Enables you to convert a volume into a replica set.For example, you can demote a volume to a failback replica set as part of a failback operation.

About Permanently Promoting a Replica Set to a Volume

You can permanently promote a replica set in a single operation, resulting in a new standard volume, template volume, or thin clone volume. You might need to perform this task if the original volume is destroyed. Permanently promoting an inbound replica set does not require any additional space on the secondary group, because it reduces delegated space by the size of the volume’s replica reserve.

Prerequisites for Permanently Promoting a Replica Set to a Volume

The following constraints apply:

- In some cases, you cannot permanently promote a replica set in a single operation. If you cannot deselect the Keep ability to demote to replica set option, you must temporarily promote the replica set and then make the promotion permanent. See Promote an Inbound Replica Set to a Recovery Volume and Make an Inbound Replica Set Promotion Permanent.
- Before you can permanently promote a template replica set, you must permanently promote all the attached thin clone replica sets.

You need the following information:

- A new volume name must be unique in the group. Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- Access controls for the recovery volume. You can specify one or more CHAP user names, an IP address, or an iSCSI initiator name.
- Whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume. See Allow or Disallow Multihost Volume Access.
- Before permanently promoting a replica set, you must stop all pending failback operations on the volume or replications to the primary group.

Permanently Promote a Replica Set to a Volume

To permanently promote a replica set to a volume in one operation:

1. Log in to the secondary group.
2. Click **Replication**.
3. Expand the partner and then expand **Inbound Replicas**.
4. Select the replica set name.
5. Click **Promote to volume**.
6. Confirm that you want to pause inbound replication from the partner to open the **Promote replica set – Volume options** dialog box. In the dialog box:
   - Choose whether to set the volume online or offline.
   - Choose whether to retain the iSCSI target name of the original volume. (Using the target name can facilitate initiator access to the volume.)
   - Clear the **Keep ability to demote to replica set** option.
7. Click **Next** to open the Promote Replica Set – Volume Settings dialog box:
   - Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
   - (Optional) Type a description.
   - Select the storage pool.
8. Click **Next** to open the Promote Replica Set – iSCSI Access dialog box and specify the following information:
   - Select what kind of access type to apply to the volume (Copy another volume's controls, Define access control policy, or one or more basic access points). You can also choose to not allow access.
   - Select the access controls using the options you see based on the access type you chose.
   - Choose whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume.
9. Click **Next**.
10. Review the information in the **Promote replica set – Summary** dialog box.
11. Click **Finish**, or click **Back** to make changes.

The replica set disappears from the list of inbound replicas, and the new volume appears in the list of volumes.

About Failing Over and Failing Back a Volume

If a failure or maintenance in the primary group makes a volume unavailable, you can fail over to the secondary group and allow users to access the volume. If the primary group becomes available, you can fail back to the primary group.

Restriction: You cannot replicate a recovery template volume, and you cannot demote a template volume to a failback replica set.

After you complete the operation, initiators can connect to the volume on the primary group and replication can continue as usual. By default, the failback baseline is reestablished.

NOTE: Snapshots must be offline for failback to occur.

Example of Failing Over and Failing Back a Volume

An example of how to fail over a volume to the secondary group and then fail back to the primary group is shown in Figure 21. No Failure (Data Available) to Figure 25. Step 3–Fail Back to the Primary Group.

Figure 21. No Failure (Data Available) shows the replication configuration, where GroupA is replicating Volume1 to GroupB.

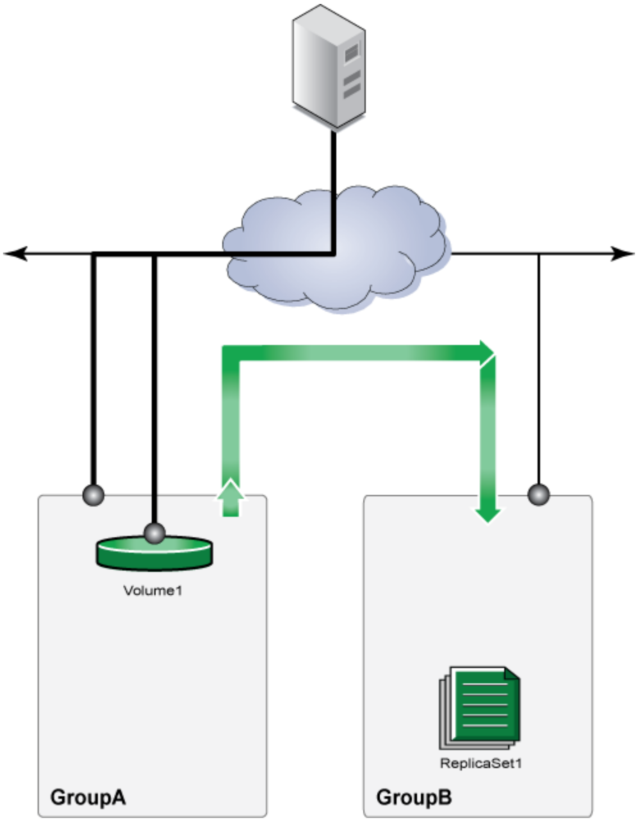

Figure 21. No Failure (Data Available)

Figure 22. Primary Group Failure (Data Not Available) shows the replication configuration after a failure in the primary group (GroupA).

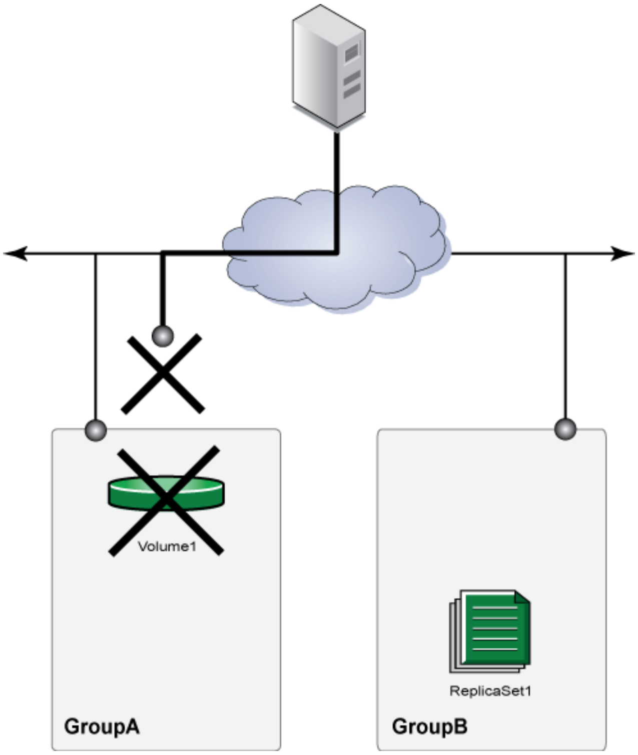

Figure 22. Primary Group Failure (Data Not Available)

Figure 23. Step 1–Fail Over to the Secondary Group (Data Available) shows the first step in recovering data on the secondary group, which is to fail over the volume to the secondary group. To fail over the volume, promote the inbound replica set to a recovery volume and snapshots. The recovery volume contains the volume data represented by the most recent, complete replica. Users can connect to the recovery volume to resume access to volume data.

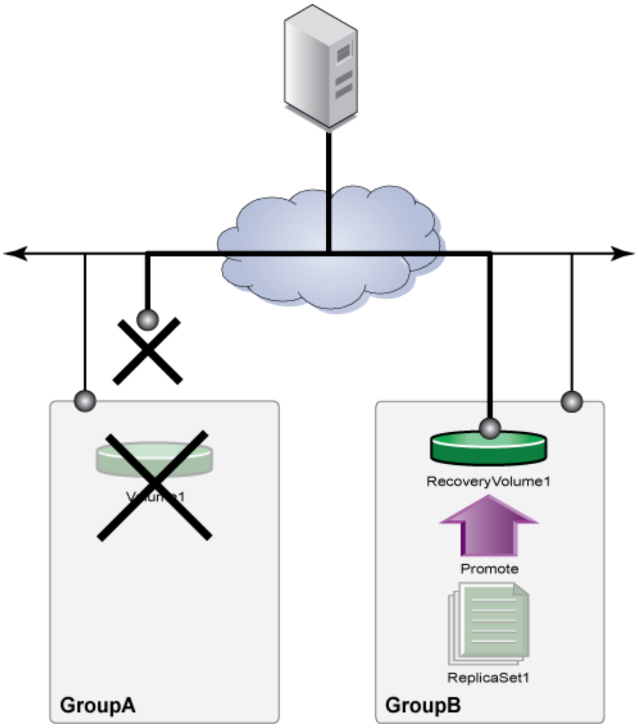

Figure 23. Step 1–Fail Over to the Secondary Group (Data Available)

Figure 24. Step 2–Replicate to the Primary Group (Data Available and Protected) shows the second step in recovering data—replicate to the primary group. When the primary group is available:

- Demote the original volume to a failback replica set.
- Replicate the recovery volume to the primary group.

NOTE: If the failback snapshot is not available on the primary group, the first replication transfers all the recovery volume data, instead of only the changes that users made to the recovery volume.

You can perform these tasks separately or use the Replicate to Primary operation, which encompasses both tasks.

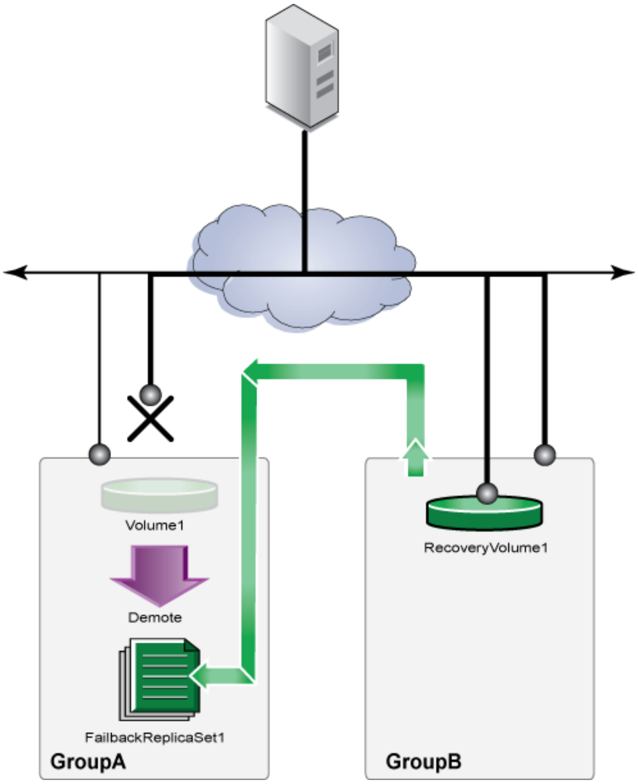

Figure 24. Step 2–Replicate to the Primary Group (Data Available and Protected)

Figure 25. Step 3–Fail Back to the Primary Group shows the final step in recovering data — fail back to the primary group. To fail back to the primary group:

- Set the recovery volume offline.
- Replicate the recovery volume to synchronize volume data across both groups.
- Demote the recovery volume to an inbound replica set.
- Promote the failback replica set to a volume.

You can perform these tasks separately or use the Failback to Primary operation, which encompasses all the tasks.

At this point, the volume replication configuration returns to its original state and users can connect to the volume on the primary group.

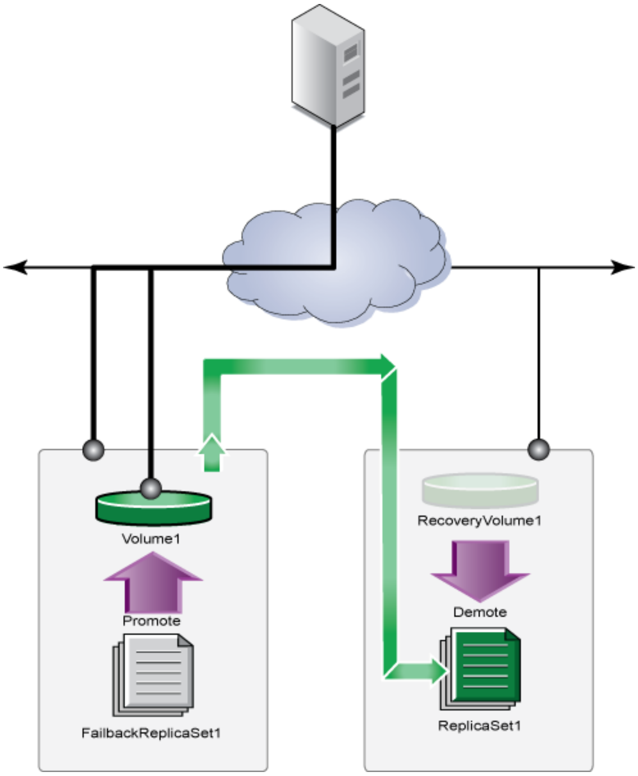

Figure 25. Step 3–Fail Back to the Primary Group

About NAS Disaster Recovery

Disaster recovery restores data on a primary storage resource and returns that resource to a full working state with minimal data loss after operation on that resource is interrupted. The interruption could be planned, such as a maintenance update, or unplanned, such as a power outage.

⚠ **CAUTION: If the site containing the source container incurs a catastrophic loss, contact Dell Technical Support for assistance.**

Disaster recovery requires that:

- A replication partnership is established between a source resource and a replica resource.
- The replica resource contains at least one replica of the source resource.

On a NAS container, disaster recovery consists of the following operations to restore data access:

- Failing over to a recovery volume
- Replicating to a partner volume
- Failing back to a source volume

You can perform these operations manually, or you can perform the replicate to partner and fail back to a primary volume operations through a single-step process.

About NAS Replication

NAS replication enables a system administrator to have an up-to-date backup system that is ready to go live if the source cluster goes offline. Similar to block volume replication, NAS container replication ensures the availability of cluster data in the event of a disaster or the NAS cluster becoming otherwise inaccessible to hosts.

NAS replication uses the Dell Fluid File System (FluidFS) snapshot-based replication technology to copy data from a source resource in the source NAS cluster to a destination resource (replica resource) in the destination NAS cluster. During replication, the system opens several TCP ports to mirror differences across the network. Figure 26. Basic NAS Replication shows an example of basic NAS replication.

When replication finishes, the system creates a replication snapshot and compares the replication snapshot on the destination NAS cluster to the replication snapshot on the source NAS cluster. Data flows in both directions in NAS replication, meaning that the same cluster can host both source and destination clusters. The system replicates only incremental changes, which improves network bandwidth utilization. The data is always consistent on the partner site and available as read-only.

NOTE: You must have group administrator privileges to perform NAS replication operations.

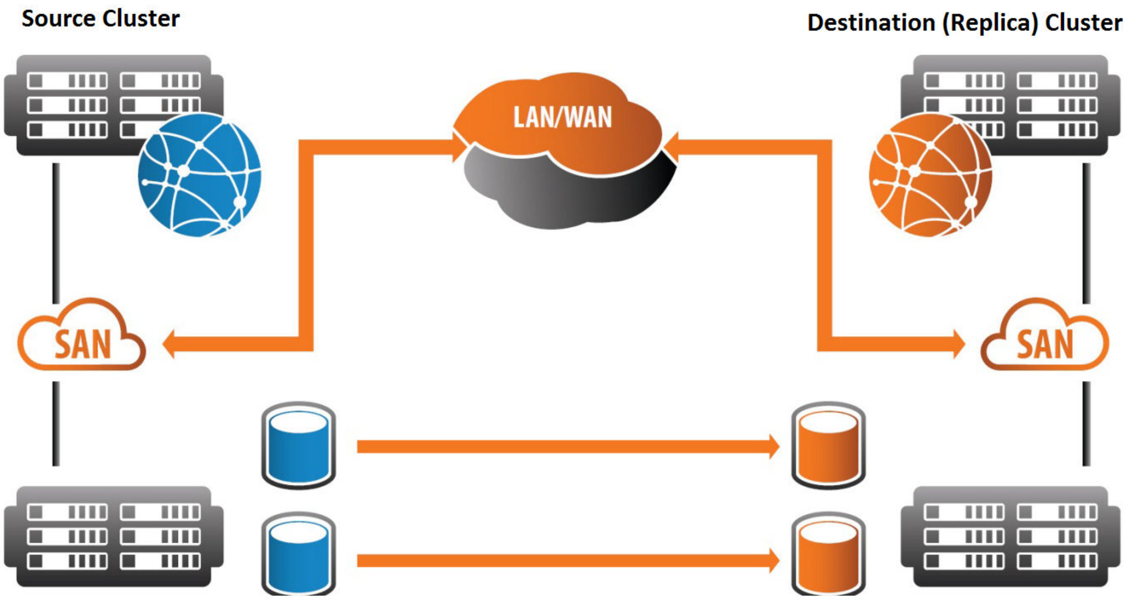

Figure 26. Basic NAS Replication

NAS volume replication consists of the following phases:

- Configure NAS volume replication — Create NAS volume replication partnerships, configure NAS volumes for replication, and delete replication relationships.
- Replicate NAS clusters or volumes — Perform manual and scheduled replication operations on NAS volumes, pause and resume NAS container replication.
- Recover NAS volumes — Promote replica volumes to recovery volumes, fail over to recovery volumes, recover data to the source volume, and fail back to the source volume. Data recovery and failback to a source volume can be performed manually or automatically.

When you perform replication on your system as part of recovering from a disaster, the type of replication you perform is determined by the state of your storage resources.

- Manual replication — Use this process if you are not able to restore data to the primary volume after failing over to the recovery volume. For example, if the primary volume must be replaced, you must perform manual replication.
- Single-step replication failback — Use this process when the primary volume is available after failing over to the recovery volume. For example, if the primary volume was taken offline for an update, you can perform a single-step replication failback.

Manually Replicate a NAS Volume

Before you can replicate a volume, the volume must be configured for replication and must have a replication partner configured.

The amount of time that the replication takes to complete depends on the size of the NAS volume and the network connection.

To replicate a NAS volume:

1. On the source cluster, select NAS, expand Local Volumes, and select the NAS volume that you are replicating.
2. Click **Replicate**.
3. Click **Yes**.
4. (Optional) Display the Alarms and Operations toolbar and click the **Failback Operations** tab.

NAS Replication Network Architecture

NAS replication depends on specific network capabilities, such as TCP ports opened over a secure tunnel and communication using the storage area network (SAN).

NAS replication opens several TCP ports. The ports are opened over a secure tunnel between the source NAS cluster and the destination NAS cluster. The ports are opened during each replication cycle.

Communication between clusters occurs over several ports as well as in both directions over the SAN.

Data flow for NAS replication occurs over the SAN network exclusively. Data flows between Ethernet ports on each NAS controller. See Figure 27. NAS Replication Data Flow for a high-level diagram of NAS replication data flow.

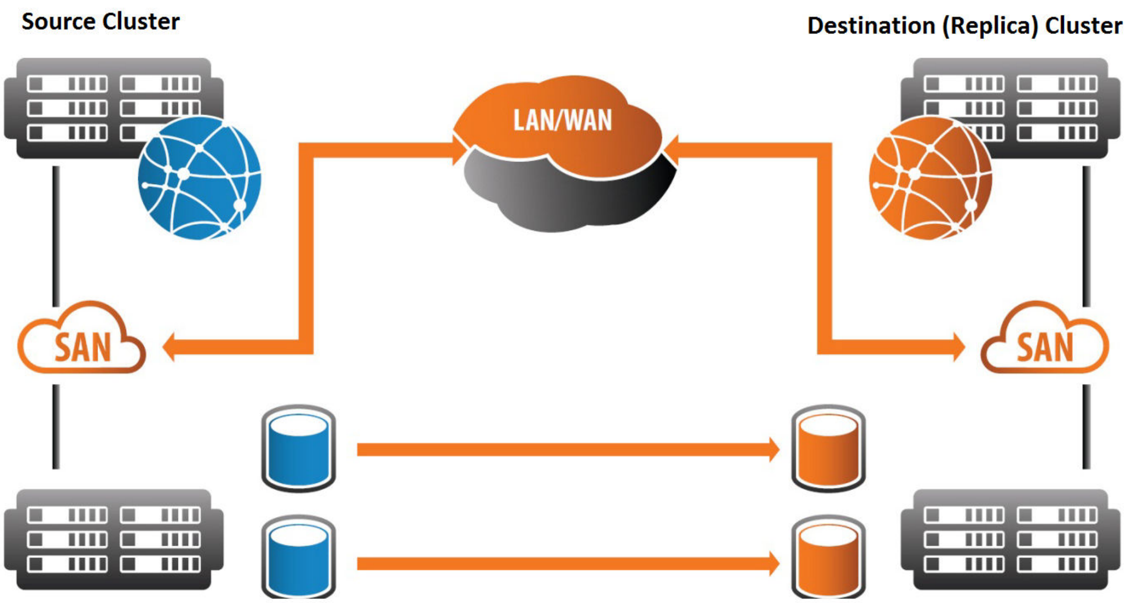

Figure 27. NAS Replication Data Flow

Network Setup and Component Requirements

Specific network configuration requirements must be met for replication to function.

NAS replication network requirements include:

- A path through the firewall for the SAN's physical and logical ports/IP addresses used on the:
  - Dell EqualLogic arrays
  - NAS controllers
- A path through the firewall for the appropriate TCP/UDP ports
- Static routes (required only when using a wide area network [WAN])

NOTE: As of v3.0, jumbo frames are not required for WANs or the SAN network.

The following physical NIC (Network Interface Controller) ports and IP addresses must be accessible through the firewall.

SAN ports:

- All EqualLogic SAN ports by way of the EqualLogic Group IP
- NAS cluster SAN Management Virtual IP (VIP) address
- Physical SAN ports on every NAS controller and SAN IP address

NOTE: Ports referred to here are physical and do not refer to TCP/IP port numbers opened through the TCP/IP stack using an application.

Table 51. TCP/IP Port Numbers shows the ports that must be open on the firewall.

Table 51. TCP/IP Port Numbers

Function	Protocol	Port Number(s)
Replication setup ports	SSH, TCP	22, 26, and 3260
Replication data ports	TCP	2-Controller Cluster: 10560-10568 4-Controller Cluster: 10560-10576
Ephemeral ports	TCP	40000 - 65535

NOTE: The ephemeral port range defines the range of TCP ports allocated by the TCP/IP stack on EqualLogic arrays and FS Series NAS appliances for initiating communications to other EqualLogic arrays or FS Series NAS appliances. For example, when configuring replication between two NAS clusters, an ephemeral port (from the range in table above), such as 41020, is selected on the source NAS cluster and establishes a connection to TCP port 22 on the destination NAS cluster.

NAS Replication WAN Network Requirements

Replication is possible over a WAN only when the partner system is accessible through a gateway to the Internet or private WAN. When replicating over a WAN, you must have a static route in addition to the regular partnership setup. The WAN can be the public Internet or a private network. Because NAS replication transfers data over the SAN network, you have to have static routes to route replication packets through the designated gateway on the SAN network.

A static route provides a network path for the NAS cluster and the NAS controllers to the SAN IP addresses for the partnered NAS cluster and NAS controllers.

When you configure NAS replication between two clusters, static routes are added on both clusters automatically from the list of default gateways on members in the EQL group. Optionally, you can change the automatically specified static routes. You can specify a different gateway for the static route or add or delete static routes for replication if the automatic method fails for any reason. The number of static routes that can be configured is not limited. If a static route needs to be modified, delete the static route and then add the static route back with the updated information.

NOTE: You can manually configure the NAS static routes only using the CLI. You cannot configure static routes using the Group Manager GUI.

If you modify the SAN network settings of a cluster, the software tries to automatically update all the partner clusters. However, to ensure continued operation, static routes on all the partner clusters should be verified by the administrators.

NOTE:

- **The static routes need to be updated on partner clusters, not the cluster whose SAN network changed.**
- **Jumbo frames are not required when connecting over a WAN network.**

You can use the following SAN network CLI commands to add, delete, or show static routes:

- nas-cluster select network select static-route add
- nas-cluster select network select static-route delete
- nas-cluster select network select static-route show

For more information about these CLI commands, see the *Dell EqualLogic Group Manager CLI Reference Guide*.

FS7610 Cluster Management and Replication Port

CAUTION: The ports identified with an arrow in Figure 28. FS7610 NAS Cluster Management and Replication Port are used for NAS cluster management and replication functionality. It is critical that these ports remain connected and operational at all times.

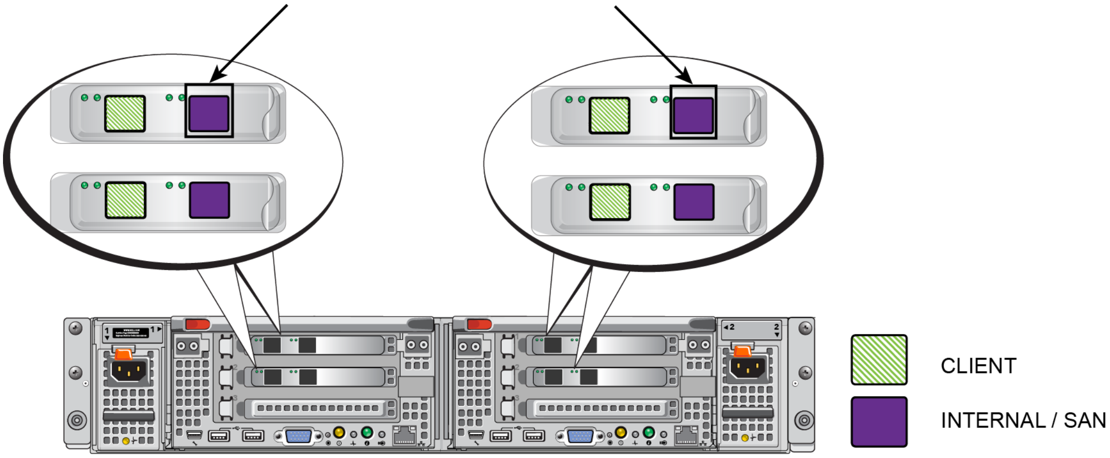

Figure 28. FS7610 NAS Cluster Management and Replication Port

FS7600 Cluster Management and Replication Port

CAUTION: The ports identified with an arrow in Figure 29. FS7600 NAS Cluster Management and Replication Portare used for NAS cluster management and replication functionality. It is critical that these ports remain connected and operational at all times.

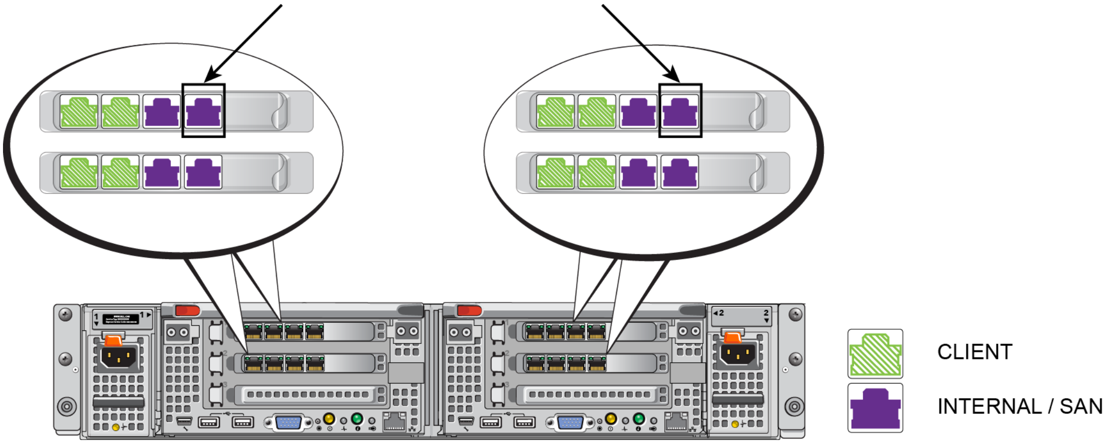

Figure 29. FS7600 NAS Cluster Management and Replication Port

FS7500 Cluster Management and Replication Port

CAUTION: The port identified with an arrow in Figure 30. FS7500 NAS Cluster Management and Replication Port is used for NAS cluster management and replication functionality. It is critical that this port remain connected and operational at all times.

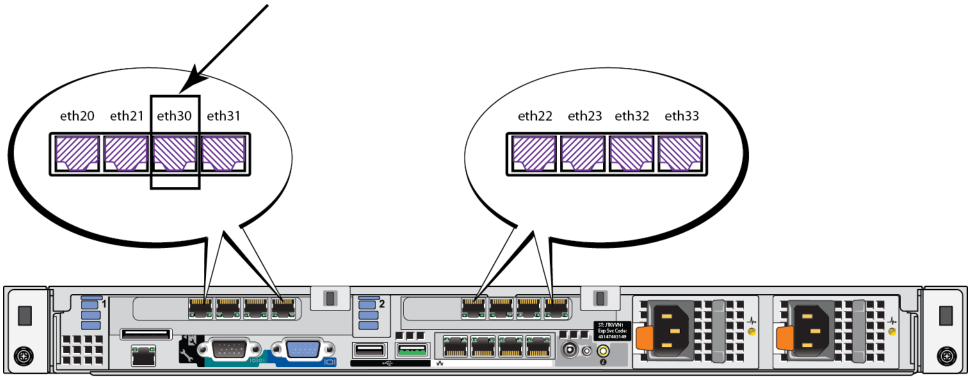

Figure 30. FS7500 NAS Cluster Management and Replication Port

Set Up Your NAS Replication Environment

To help ensure successful replication, for each NAS container that you want to replicate, follow these steps to set up your replication environment:

1. Gather the following information to help you determine how much replication space you need:
   - Number of replicas that you want to keep
   - Average time span between each consecutive replica
   - Reported size of the volume
   - Whether thin-provisioned
2. Identify a replication partner (secondary group) to store the volume replicas. This secondary group must meet the space and network connectivity requirements.
3. Verify on the primary group that the secondary group is configured as a replication partner.
4. Verify on the secondary group that the primary group is configured as a replication partner and that free space has been delegated to hold the replicas.
5. On the primary group, configure the volume for replication, specify the appropriate replication space values, and verify the replication partnership.
6. (Optional) Set up a schedule to create replicas on a regular basis.
7. Monitor each replication operation and make sure it is successful.

   If the replication operation is not successful, identify and correct the problem. For example, you might need to increase network bandwidth or increase replication space.

NAS Volume Replication Operations

The following operations are part of NAS container replication:

- Configure NAS Container Replication Partners

  As with volume replication, NAS container replication requires one or more replication partners, with their own configuration requirements. This replication relationship can be tested at any time after you configure the relationship.
- Configure Containers for Replication

  Select the source container and specify a remote (destination) replica container on the replication partner.
- Replicate NAS Containers

  In the default configuration, NAS container replication has no predefined schedule and runs on demand. If the source and destination clusters can communicate with each other, you can replicate a NAS container configured for replication at any time.
- Schedule NAS Volume Replication

  NAS volume replication operations can be scheduled to run at a specified date and time, or at fixed hourly or daily intervals.
- Pause NAS Container Replication

Pausing NAS container replication suspends any replication operations for the container that are in process. While replication is paused, scheduled replications do not take place.

You can pause NAS container replication for individual containers. Unlike volume replication, it cannot be paused for all replications to a specific partner. You can pause replication from either the source or destination group.

- Resume NAS Container Replication

  If the source and destination groups can communicate with each other, you can resume NAS container replication at any time. When you resume the replication, any replication operations that were in progress at the time the operation was paused will resume, and any NAS container replication schedules will resume at their next scheduled time. You can resume replication from either the source or destination group.

- Cancel NAS Container Replication

  When you cancel NAS container replication, replication operations that are running on the specified container stop and the container's replica remains in the state it was in before the canceled replication began.

  You can cancel replication for individual NAS containers, but you cannot cancel all in-progress replication operations taking place in a NAS cluster or between a specified pair of replication partners. Replication operations can be canceled only from the source cluster.

- Delete a NAS Container Replication Relationship

  When you delete replication for a NAS container, the replication relationship between the source and destination containers is discontinued and the replica container is deleted. If the replica has been promoted to a recovery container, the recovery container is not deleted.

  You can delete replication for individual containers, but you cannot simultaneously delete replication for all containers in a cluster, or all container replication pairings with a specific partner. Deleting replication for a NAS container, like disabling replication for a volume, does not disrupt other replication operations or disrupt the replication partnership between the source and destination groups. NAS container replication can be deleted from either the source or destination cluster.

  The replication relationship is canceled when you delete a NAS container, but the replica container in the destination cluster remains and is automatically promoted to a standalone container residing on the destination cluster.

About NAS Container Replication Partners

As with volume replication, NAS container replication requires that you pair the group containing the source NAS cluster with a replication partner. After you establish a replication partnership, you can configure individual containers for replication.

Each cluster can have multiple replication partners. Using multiple replication partners enables you to replicate different containers to different partners, depending on the operational requirements. For example, you can replicate containers used by two different departments to two different partners. However, each individual container can be configured to replicate to one replica container on one partner only.

To change a container's replication partner, you must delete the container's replication configuration and then reconfigure replication using the new partner.

⚠ **CAUTION: Deleting a container's replication configuration also deletes that container's replicas.**

NAS replication partners are configured using the same process used when configuring volume replication partners. You can enable both NAS container replication and volume replication on the same partner, and test the IP address and authentication used with a replication partner.

NAS Container Replication Partner Requirements

NAS container replication partners must meet the following requirements:

- Source and destination NAS clusters must be running the same FS Series firmware versions.
- Source and destination NAS clusters do not have to be running the same hardware. For example, you can replicate from an FS7500 cluster to an FS7610 cluster.
- The network connection between the replication partners must be capable of supporting replication data transfers. The Manual Transfer Utility (MTU) cannot be used to facilitate NAS container replication.

NOTE:

- **The requirements for NAS container replication are not validated when you create the replication partnership. The system will not allow you to configure replication for a container if the source and destination clusters do not meet the configuration requirements.**
- **Restoring volume configuration on the destination cluster is not supported between major revisions (such as version 3 > version 4).**

Before configuring your network for replication, see the *Installation and Setup Guide* to complete the following tasks:

- Allocate and configure the NAS cluster
- Configure and assign an IP address for all network ports
- Verify that the correct ports are connected and remain operational at all times on the NAS systems

About Replicating to a Partner

Replicating to a partner uses the partnership configured in the failover process to recover data to the source, or primary, container. After the failover, data on the recovery container is written to the primary container. If you are replicating to a new container, all of the data on the recovery container is written to the newly configured primary container.

If the original primary container is no longer available and the recovery container is configured to replicate to another container, the replication process takes longer than it would if you were replicating back to the original primary container. Replicating to the original source container writes only data written to the recovery container after the failover back to the source container. However, replicating to a new source container writes all of the data to the new source container, and might affect client access performance.

You can perform these operations manually as part of the replication process or you can perform them as part of the single-step failback process. Time and bandwidth, as well as the amount of data being written to the source container, affect how much time the single-step failback process takes.

Manually Replicate a NAS Container

Before you can replicate a container, the container must be configured for replication and must have a replication partner configured.

The amount of time that the replication takes to complete depends on the size of the NAS container and the network connection.

To replicate a NAS container:

1. On the source cluster, click **NAS**, expand **Local Containers**, and select the NAS container that you are replicating.
2. Click **Replicate**.
3. Click **Yes**.
4. (Optional) Display the Alarms and Operations toolbar and click the **Failback Operations** tab.

Pause or Resume Replication of a NAS Container

1. In the navigation menu, click **NAS**, then select the NAS cluster on which the container resides.
2. Expand **Local Containers** and select the container.
3. In the Activities panel, either:
   - Click **Pause replication** to pause replication.
   - Click **Resume replication** to resume replication.
4. Click **OK**.

Activate the Destination Cluster

After you promote the destination containers, configure the destination cluster to serve client requests:

- Apply the same cluster settings to the destination cluster that were configured on the source cluster.
- Enable the destination cluster to perform the same function as the failed source cluster.

The configuration of your environment determines how you change the configuration. For example, if the source cluster uses Active Directory (AD) / Lightweight Directory Access Protocol (LDAP), the destination cluster must use the same AD/LDAP. This setup ensures all user information is retained in the new configuration.

- Change the Domain Name System (DNS) server to point to the destination cluster instead of the source cluster. You should make the change on the DNS server used to resolve NAS cluster virtual IPs (VIPs), if any.
- If the NAS cluster has not already been joined to AD or LDAP/NIS, connect them now.

During the cluster activation period, client connections might fail and need to be reestablished. However, the change should be transparent to the client as long as they continue to use the same fully qualified domain name (FQDN) to reach the Server Message Block (SMB) share or Network File System (NFS) export.

NOTE: Restoring volume configuration on the destination cluster is not supported between major revisions (such as version 2 > version 3).

About Single-Step Failback to Primary

You can use the single-step failback to primary process as part of the disaster-recovery process.

To prepare for disaster recovery, a primary (source) container is configured to replicate data to a replica container. Clients have read-write access only to the source container. When the source container becomes unavailable, you can perform a *failover* by temporarily promoting the replica container, defining a new replication partnership between the source container and the recovery container, and giving clients read-write permissions to the recovery container.

The last version of data replicated to the replica container (now the recovery container) is written to the source container. When the source container is available again after the failover, you perform a *failback*, which writes to the source container any additional changes made by the clients on the recovery container while the failover was in place, and demotes the recovery container to a replica container. The replication partnership is redefined between the source container and the replica container, and clients are given read-write permissions to the source container again.

For example: Container A is a source container with which clients interact directly. Container A has a replication partnership with container B, a replica container. If container A must be taken offline for maintenance, container B is temporarily promoted to a recovery container, a new replication partnership is defined between container A and container B, and clients interact directly only with container B. When container A is available after maintenance, all of the latest updates from clients are replicated from container B to container A, then container B is demoted to a replica container and clients resume interacting directly only with container A.

If the source container is no longer available, establish a replication relationship between the replica container and a new source container, manually perform replication from the replica container to the new source container, then demote the recovery container to a replica container.

The single-step failback process performs the replication and failback operations with minimal user interaction. When you perform a single-step failback operation (after failing over) without performing a manual replication, the following actions occur:

1. Data is replicated from the recovery container to the source container.
2. The recovery container is demoted to a replica container.
3. A replication partnership is redefined between the source container and the demoted replica container, and the source container resumes replicating to the replica container.

   CAUTION: To ensure that no data is lost, verify that clients cannot write to the recovery container during the single-step failback process.

Perform Single-Step Failback to Primary

Single-step failback to primary performs data replication and failback with minimal user interaction.

Before you can perform this task:

- You must be logged in as group administrator on the group that contains the recovery container.
- You must know the Group Manager account user name and password for the group that contains the source container.
- You must have temporarily promoted the replica container to a recovery container.

- Both groups must be running PS Series firmware version 7.0 or later, and the clusters on those groups must be running FS Series firmware version 3.0 or later.

To perform single-step failback to primary:

1. Click **NAS**, expand **NAS Cluster**, and expand **Local Containers**.
2. Select the recovery container.
3. Click **Failback to primary**. The Failback to Primary message is displayed.
4. Click **Yes**. The Replicate Recovery Container message is displayed.
5. Type the account name and password for the array on the primary site and click **OK**.

   To monitor the failback operation on the recovery container, display the Alarms and Operations toolbar and click the **Failback Operations** tab.

About NAS Replica Containers

NAS replica containers reside on the remote (destination) replication partner cluster and contain replicated data from the source container. As with replica volumes, you can temporarily or permanently promote replica containers to recovery containers and grant hosts access to recovery container data.

Unlike source containers, you cannot create snapshots of replica containers. However, any replicas of the source container are replicated to the destination partner.

The following considerations apply to replica containers:

- The NAS container reserve on the replication partner must have enough free space to store the replica.
- Unlike volume replication, which retains a replica set of previously created replicas, the system retains only the current replica of the source container in a NAS container replication configuration. To roll back to a previous point in time, you must use snapshots.
- You can either replicate the container to a new replica container or you can specify that the container be replicated to an existing replica container.
- The system assigns names to new replica containers as follows:
  - The group gives the replica container the name of the source container, with a randomly generated four-digit string appended to the end of its name. For example, if you replicate a container named `ContainerA`, the replica might be named `ContainerA_8970`.
  - If the length of the container name, combined with that of the appended string, is more than 229 characters, the system shortens the replica container name to be 229 characters long. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.
- If you replicate to an existing container:
  - The destination container must either be an empty container containing no data, or a container to which the source container has previously replicated.
  - The source and replica containers do not need to be the same size. Whenever replication occurs, the group automatically adjusts the size of the replica container so that it either shrinks or grows to match that of the source container.
  - All data previously residing on the container is overwritten and cannot be recovered.
- If read-only access is enabled on a replica container, it appears in the list of containers on the destination cluster.

  NOTE: Replica containers count toward the total number of containers in the group. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for the maximum number of containers allowed in a group.
- Container Replication of configuration information about SMB shares and NFS exports from a NAS container to a replica container might be delayed for several minutes. This delay will have no effect on the replication of the data itself, but might affect your ability to access the data because you might have to recreate recent configuration updates. If you replicate immediately, the first replication might not have all the data, but subsequent replications will.

Replicate to Another Container in a Cluster

If the original source container is no longer available, you can configure the recovery container to replicate to another container in the original source cluster. However, if the original source container is available, Dell recommends failing back to it. Failing back to the original source container takes less time than failing back to a new container.

To replicate to a container in a cluster:

1. Log in to Group Manager and delete replication for the source container.
2. From the destination cluster, configure replication for the promoted recovery container, specifying that it replicate back to the original source container.
3. Manually perform replication on the promoted recovery container.
4. After replication completes, log in to the source cluster and promote the original source container.
5. Log in to the original destination cluster and delete replication for the promoted recovery container.
6. Log in to the original source cluster and configure replication between the source container and the promoted recovery container.

   NOTE: Restoring volume configuration on the destination cluster is not supported between major revisions (such as version 2 > version 3).

Configure a Container for Replication

As with volume replication, NAS container replication requires that you pair the group containing the source NAS cluster with a replication partner. After you establish a replication partnership, you can configure individual containers for replication.

To configure a NAS container for replication:

1. Log in to the group containing the NAS cluster that you are replicating.
2. In the navigation menu, click **NAS**, select the NAS cluster, and expand **Local Containers**.
3. Select the container for which you are configuring replication.
4. In the Activities panel, click **Configure replication**.
5. In the **Container Replication Partner** section of the Configure Replication dialog box, select the partner to which the container will be replicated.
6. (Optional) If you want to replicate the container to an existing container specified by name, select **Advanced** to open the Replica Container section of the dialog box.

   CAUTION: The initial replication operation will destroy all data previously residing on the container.

   NOTE: If you replicate to an existing container:

   - The destination container must either be an empty container containing no data or a container to which the source container has previously replicated.
   - The source and replica containers do not need to be the same size. Whenever replication occurs, the group automatically adjusts the size of the replica container so that it either shrinks or grows to match that of the source container.
7. Click **OK**.

Monitor NAS Inbound or Outbound Replication

You can monitor all NAS containers that are replicated to a PS Series group from a partner group.

1. Click **Monitoring** in the navigation menu.
2. Below NAS Replication, select either **Inbound NAS Replication** or **Outbound NAS Replication**.

View Container Replication History

1. Log in to the source cluster.
2. Click **NAS** and select the NAS cluster on which the container resides.
3. Expand **Local Containers** and select the container.
4. Click the **Replication** tab.

NOTE: To display replication history:

1. Click **Monitoring**.
2. Below NAS Replication, select **Outbound Replica Containers**.
3. Select the **Replication History** button.

Activate the Source Cluster

When you configure the source cluster to serve client requests, you reverse the changes that you made when you activated the destination cluster for failover.

While the source cluster is being activated, client connections might fail and need to be reestablished. However, the change should be transparent to clients as long as they continue to use the same FQDN to reach the SMB share or NFS export.

Your environment will determine how you change the configuration.

To activate the source cluster, either:

- Change the Domain Name System (DNS) server to point to the source cluster. Make the change on the DNS server used to resolve any NAS cluster virtual IPs (VIPs).
- Join the NAS cluster to AD or LDAP/NIS, if it has not been done already.

Create a NAS Container Replication Schedule

Before you can define a replication schedule, you must:

- Configure replication partners
- Configure replication between two NAS containers on the partners

To create a NAS container replication schedule:

1. Log in to the source cluster.
2. Click **NAS**, expand **Local Containers**, and then select a NAS container on which you have configured replication.
3. Click **Create schedule**. The Schedule Settings tab of the Create Schedule wizard opens.
4. In **Name**, type a name for the replication schedule.
5. Select the **Replication schedule** option.
6. Specify one of the following options and click **Next**:
   - A frequency at which the schedule should run: once, hourly, daily, schedule
   - To use an existing schedule as a template for the new schedule
7. If you are using an existing schedule, in the Select Existing Schedule tab, select the schedule to use as a template for the new schedule and click **Next**.
8. On the Schedule Type tab, specify the following information, as appropriate, and click **Next**:
   - Start date
   - End date
   - Time to start running the schedule
   - Repeat interval
9. On the Summary tab, either:
   - Click **Copy** to copy the schedule summary, which you can paste in a text editor such as Notepad.
   - Click **Back** to modify any of the replication schedule settings.
   - Click **Finish** to create the replication schedule.

Configure a Replication Partner

For both volume and NAS replication, you must configure the replication partnership on the two replication partners.

To configure a replication partner:

1. Click **Replication → Replication Partners**.
2. In the Activities panel, click **Configure partner**. The Replication Partner Identification tab of the Configure Replication Partner wizard opens.
3. Provide the requested information in each step of the wizard and click **Next**.

   Refer to the online help resources that are available in the user interface to view descriptions of individual fields.
4. Review the configuration and either:
   - Click **Back** to change replication partner configuration settings and repeat the configuration process.
   - Click **Copy** to copy a text version of the summary partner configuration information that you can paste into a text editor.
   - Click **Finish** to complete configuring the replication partner. The Test Partnership message appears.
5. Either:
   - Click **Yes** if the replication partnership with the local group is configured on the remote partner and that partner is available. The Test Replication dialog box opens with the results of the replication partnership test.
   - Click **No** if the partner is not configured or available.
6. If you have not already done so, repeat this procedure on the replication partner specified on the Replication Partner Specification tab.

After you configure both replication partners, you can replicate data between the partners.

Test a Replication Partnership

When you test a replication partnership, you verify the IP address and authentication defined for a replication partner. You can test a replication partnership when you complete the Configure Replication Partner wizard and at any time after you have defined a replication partnership.

If you do not test the partnership, any configuration problems between the configured partners are not identified until the source first tries replicating to the replication partner.

To test a replication partnership:

1. Click **Replication** and select a replication partner.
2. Click **Test partnership**.
3. Click the **Test Partnership** button.

Test Replication Partnership Results

Table 52. Test Replication Partnership Results

Attribute Tested	Status	Action
Partner IP Address	OK	Both of the following conditions have been verified: • The configured partner's IP address can be reached. • A partnership with the local group is configured on the remote group.
	Incorrect IP address	At least one of the following issues occurred: • The configured partner's IP address could not be reached. • The IP address does not belong to a partner's group. To resolve, click the **Modify replication partner IP address** link and verify that the IP address is configured correctly.
	Partner is not configured	The IP address can be reached, but a partnership with the local group is not configured on the remote group. To resolve, click the link to return to the wizard tab, where you can configure replication on the remote group.

Attribute Tested	Status	Action
	Test was not run	An internal error prevented this test from running. If you continue to receive this error, contact Dell customer support.
Authentication	OK	All of the following conditions have been validated: • The configured partner's name matches the remote group name. • The inbound and outbound passwords configured on the remote replication partner.
	Invalid	At least one of the following issues occurred: • The partner name does not match the configured replication partner's group name. • The local group name does not match the partner name that is configured on the remote group for the configured replication partnership. • The local outbound password does not match the inbound password configured on the replication partner. • The local inbound password does not match the outbound password configured on the replication partner. To resolve, verify the partner and group names and the configured passwords. Click the **Modify replication partner IP address** link and verify that the IP address is configured correctly.
	Test was not run	The replication partner IP address test did not complete successfully and the authentication was not tested. To resolve, click the link to return to the wizard tab, where you can resolve the problems found with the IP address verification.

Delete Replication for a NAS Container

NOTE: Deleting replication for a NAS container disables replication to the destination container and destroys all data on the replica.

1. Log in to the source cluster.
2. In the navigation menu, click **Replication**.
3. Select the replication partner containing the replica NAS container.
4. Select **NAS Replication**.
5. Below **Outbound Replicas**, select the container for which you are canceling replication.
6. In the Activities panel, click **Delete replication**.
7. Click **OK**.

NOTE: You can also delete replication from the destination group by selecting the an inbound replica and clicking Delete replication.

Recover from an Offline or Unavailable NAS

If a NAS goes offline or becomes unavailable, the source NAS cluster goes offline and the full functionality of that cluster is not available.

To recover the offline or unavailable NAS, you must follow this process:

1. Fail over to the destination cluster.
2. Resolve the problem with the source cluster.
3. Fail back to the source cluster.

For disaster recovery on NAS containers, after you fail over to the destination container, you can fail back to the primary container in a single-step process.

A properly configured system does not require a configuration restoration to perform a failover operation. However, if the source cluster configuration needs to be applied to the destination cluster, contact Dell Technical Support for assistance.

Fail Over to a Recovery Volume

Failing over to a recovery volume prepares storage resources to recover data that was replicated to a replica volume. A primary volume does not have to be available until the end of this process.

The process of performing a failover to the replica volume:

- Removes the replication partnership defined between the primary and replica volumes.
- Promotes the replica volume to a recovery volume, enabling clients to read and write directly with the recovery volume, instead of the primary volume.
- Defines a replication partnership between the recovery volume and a primary volume, enabling data to be restored from the recovery volume to a primary volume.
  - If the original primary volume is not immediately available, this partnership can be defined after the primary volume becomes available.
  - If the original primary volume is no longer available, this partnership is defined between the recovery volume and a new primary volume.
  - For example: Primary volume A is configured in a replication partnership with volume B. If volume A is no longer available, as part of the failover process, volume B can be configured with a new primary volume, volume C.

NOTE: The recovery volume's data is complete up to the point in time of the most recent successful replication.

If user and group quota rules are in place for the source volume, those quotas are applied to the recovery volume as follows:

- If an external server (such as Active Directory, LDAP, or NIS) is authenticating users and groups on the destination cluster, quotas are applied to the recovery volume.
- If only local users and groups are used, quotas are not applied.
- If local users and groups are used *and* an external server authenticates users and groups, user and group quotas are applied for the externally authenticated accounts. However, you must verify that quotas are applied correctly for the local users and groups. If they are not applied correctly, modify them so that they are correct.

Access Replica Container Data

To grant hosts access to the contents of a replica container, or fail over to the NAS replica container, select the **Apply Host Access Configuration** option. After you perform that operation, hosts that had access to the source container are granted access to the recovery container.

Fail Over to the NAS Destination Cluster

When a disaster occurs in a NAS network and the source NAS cluster becomes unavailable for some reason, you must fail over to the destination (or replica) NAS cluster, as shown in Figure 31. NAS Replication Failover.

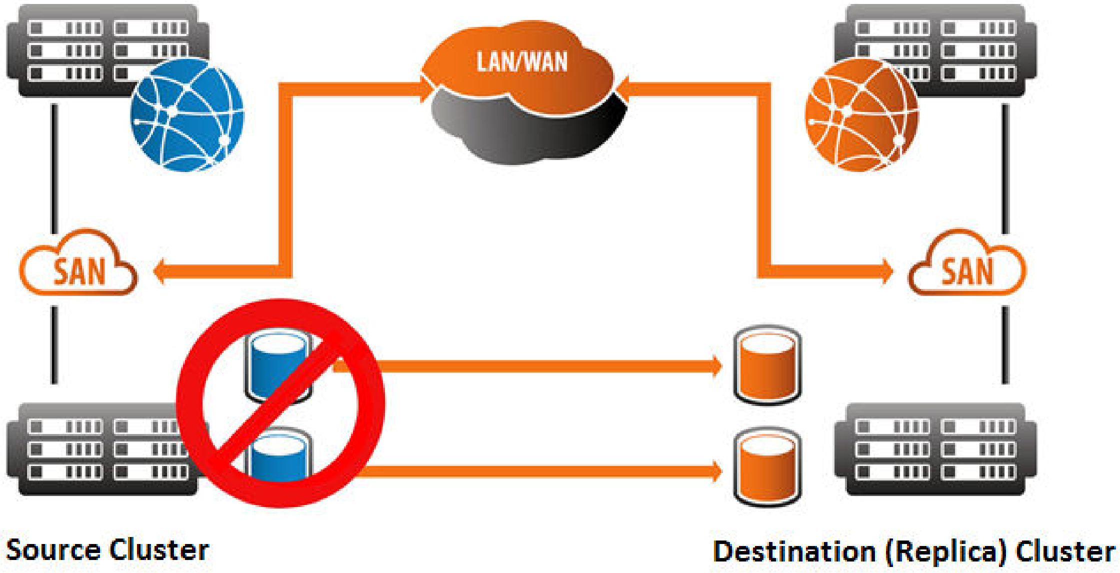

Figure 31. NAS Replication Failover

Performing a failover to the NAS destination cluster involves the following steps:

- Promoting each replication destination NAS container
- Activating the destination cluster

After resolving the cause of the failure on the NAS source cluster, fail back to the NAS source cluster.

Fail Back to a Source Volume

Failing back to the source, or primary, volume reestablishes the originally configured replication relationship.

Failing back to the source volume:

- Removes the replication partnership defined between the restored primary and recovery volumes.
- Demotes the recovery volume to a replica volume, enabling clients to read and write directly with the primary volume, instead of the replica volume.
- Defines a new replication partnership between the primary volume and the replica volume, enabling the recovery volume to restore data to the primary volume.

You can perform these operations manually as part of the failback process or you can perform them as part of the single-step failback process. Time and bandwidth, as well as the amount of data being written to the source volume, affect how much time the single-step failback process takes.

NOTE: Snapshots must be offline for failback to occur.

Fail Back to the NAS Source Cluster

After resolving the issues that led to the source NAS cluster failure, such as replacing hardware, disks, and so on, return the source and destination NAS clusters to their original functions.

To fail back to the NAS source cluster:

1. Configure the replication partnership.
2. Configure the cluster for replication.
3. Launch the replication manually.
4. Promote replica clusters.

5. Activate the source cluster.
6. Recreate the replication relationship.

NOTE: You have to reinstall Dell Fluid File System (FluidFS) on the source cluster only if the source cluster is entirely new. See the *Installation and Setup Guide* if you must reconfigure the source cluster.

About Promotions and Recovery Containers

If a NAS container becomes unavailable, you can promote a replica container to a recovery container, preserving host access to the container's data. When you promote a replica, the group converts the replica to a standalone container, called a recovery container. The data on the recovery container will be current as of the last-completed replication.

When you promote a replica, you have the option of retaining the ability to demote it at a later time, thereby making the promotion temporary. Or, alternatively, you can permanently promote it. When you permanently promote a replica, the system severs the replication relationship with the source container and the destination. With temporary promotions, the system retains the replication configuration so that it can be restored if the recovery container is later demoted.

When deciding whether to temporarily or permanently promote the replica container, use the following guidelines:

- Permanently promote the replica if the source container is unavailable or if you do not plan on restoring access to it.
- Retain the ability to demote the replica container if you plan to demote it at a later point in time.

Demotions can be performed from the destination cluster at any time, provided that the source and destination clusters can communicate with each other. The demotion operation reestablishes the replication link with the source container.

CAUTION: When you demote a recovery container, all data written to the recovery container while it was temporarily promoted will be lost.

Fail Over to a NAS Replica Container

When you fail over to a replica container, you convert the replica container to a recovery container.

To fail over to a replica NAS container:

1. Log in to Group Manager on the group in which the replica volume resides.
2. Click **Replication**, expand **Replication Partners**, and expand the partner group on which the source volume resides.
3. Expand **NAS Replication** and then expand **Inbound Replica Volumes**.
4. Select the replica volume that you want to promote to a NAS volume.
5. Click **Promote to local volume**. A confirmation message is displayed.
6. (Optional) If you want to be able to return the recovery volume to the role of replica volume after the replication, verify that the **Retain the ability to demote** checkbox is selected.
7. (Optional) Select **Retain Host Access Configuration** to refresh the list of SMB shares and NFS exports residing on the volume.

   This option is available only if **Retain the ability to demote** is selected.
8. Click **OK**.

Demote a Recovery Container

Demoting a recovery container to a replica container removes read-write access from the recovery container and reestablishes replication with the source container.

CAUTION: When you demote a recovery container to a replica, any data written to the container while it was a recovery container is destroyed.

1. Log in to the destination cluster.
2. Click **NAS**, then select the cluster containing the recovery container.
3. Expand **Local Containers** and select the recovery container.
4. In the Activities panel, click **Demote to replica container**.
5. Click **OK**.

About Cloning Volumes

Cloning a volume creates a new standard volume, template volume, or thin clone volume, with a new name and iSCSI target, but the same reported size, pool, and contents as the original volume at the time of the cloning.

- [Parent volume] Templates are read-only (gold image) copies of a volume.
- Thin clones are duplicate volumes that share space with their parent volume.

Cloning a volume consumes space from the pool where the original volume resides. The space required for cloning a volume is equal to the volume reserve at the time of the clone operation. Reserving snapshot space for the new volume requires additional pool space. Initially, you must create the cloned volume in the same pool as the original volume. After creating the cloned volume, you can move it to another pool with adequate space. If the original volume is a thin-provisioned volume, a clone of that volume will also be thin-provisioned and will initially report the same size and space utilization as the original volume.

Common use cases for cloning include:

- Reboot and restore

  You can create a "boot" master volume (template volume) and then create a (thin clone) volume for every client machine. In this scenario, you can wipe data from a computer and then restore it from a master image.
- Provision new computers

  You can provision computers with a standard set of software, minimizing time required to install individual applications.
- Full system backup

  You can create a backup of the operating system and installed software.
- Transfer to another user

  You can reset a computer to the master image prior to reallocating the computer to another user.

Clone a Volume

To clone a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the name of the volume that you want to clone.
3. Click **Clone** in the Activities panel. The Clone Volume wizard opens.
4. Specify the required information in each step and click **Next**.

   Refer to the online help resources that are available in the user interface to view descriptions of individual fields.
5. When you reach the final step, click **Finish** to create the clone.

About Cloning an Inbound Replica

Cloning lets you access the data in a replica without any impact on the replication configuration or the replica.

You can clone an inbound replica to create a new volume on the secondary group. After the clone operation completes, the replica is still available, and replication continues as usual.

The new volume has a different iSCSI target name, but contains the same data as the replica from which it was cloned. The volume will be available to iSCSI initiators immediately through the IP address of the secondary group (but not the primary group). You can access and modify the new volume as you would any other volume of its type. The new volume is located in the same pool as the replica.

Cloning different replica types leads to different results:

- Cloning an inbound replica creates a thin-provisioned volume, regardless of whether the original volume was thin-provisioned.
- Cloning a template replica creates a new template volume.

- Cloning a thin clone replica creates a new thin clone volume, which remains attached to the thin clone replica set.

By default, the new volume is set online, has read-write permissions, and uses the group default snapshot space and iSCSI settings.

Cloning a replica consumes 100 percent of the original volume reserve from free secondary group pool space. If you want to create additional snapshots or replicas of the new volume, additional space is needed.

Clone an Inbound Replica

Before you clone an inbound replica, the following considerations and decisions might apply:

- You can configure access controls.
- You can allow initiators with different iSCSI qualified names (IQNs) access to the volume and its snapshots.
- If you clone a thin clone replica to create a new thin clone volume, the new volume is attached to the template replica set on which the thin clone replica depends.
- Name can be up to 63 bytes and is case-insensitive. You can use any printable Unicode character except for ! " # $ % & ' ( ) * + , / ; < = > ?@ [\ ] ^ _ ` { | } ~. First and last characters cannot be a period, hyphen, or colon. Fewer characters are accepted for this field if you type the value as a Unicode character string, which takes up a variable number of bytes, depending on the specific character.

To clone an inbound replica:

1. Click **Replication** and then expand the replication partner.
2. Expand **Inbound Replicas** and then select the replica set.
3. In the **Replicas** panel, select the replica's timestamp.
4. In the Activities panel, under the selected replica, click **Clone replica** to open the Clone Volume Replica – Volume Settings dialog box.
5. Provide the requested information in each step of the wizard and click **Next**.

   Refer to the online help resources that are available in the user interface to view descriptions of individual fields.
6. Click **Finish**, or click **Back** to make changes.

The new volume appears in the list of volumes in the far-left panel.

18

About Synchronous Replication

Synchronous replication (SyncRep) is the simultaneous writing of data to two pools for a volume in the same PS Series group, resulting in two hardware-independent copies of the volume data. Each write must go to both pools before the write is acknowledged as complete. If one copy of the volume data is not available due to a power failure or resource outage, you can still obtain the data from the other pool.

NOTE:

- **You cannot perform traditional replication on a synchronous replication-enabled volume.**
- **To support up to 32 volumes for synchronous replication, a PS Series group must contain only PS6xx0 arrays. If the group contains one or more PS4xx0 arrays, only 4 volumes are supported.**

How Synchronous Replication Works

Synchronous replication (SyncRep) is enabled on a per-volume basis. For volumes for which synchronous replication is not enabled, volume data and snapshots are located only in the pool to which the volume is assigned. For synchronous replication-enabled volumes, volume data exists simultaneously on two copies of the volume:

- SyncActive — The active copy of the volume to which iSCSI initiators connect when reading and writing data.
- SyncAlternate — When data is written to the SyncActive volume, the group simultaneously writes the same data to this copy of the volume in the alternate pool.

You can switch a volume between being the SyncActive volume and the SyncAlternate volume. The former SyncActive volume becomes the SyncAlternate volume and vice versa. However, note that if you have more than one volume, but switch only one of them, only that volume switches from one state to the other. The other volumes retain their current status.

When you switch a volume, no iSCSI target configuration changes are required. During the switch, host connections are logged out; iSCSI initiators can reconnect when the switch has completed. Depending on its configuration, the initiator might reconnect automatically.

When synchronous replication is first enabled, or at any other time when data is being written to both volumes to become in sync, performance degradation might occur. This effect increases with the quantity of tracked changes, but it is significantly reduced after the volumes become in sync. Also, depending on the quantity of tracked changes, activity within the group, and available network bandwidth, an extended period of time might elapse before the two volumes become in sync again. The Group Manager GUI displays the status of this operation.

NOTE: If you delete a volume for which synchronous replication is enabled, the group will place the SyncActive volume into the recovery bin. However, the SyncAlternate volume will be deleted and cannot be recovered.

Compare SyncRep and Traditional Replication

The decision to use synchronous replication (SyncRep) or traditional replication should be driven primarily by your business requirement and recovery point objectives (RPO). However, you also should examine technical considerations such as networking, available capacity, and data recovery times when deciding whether traditional replication or synchronous replication is best suited for protecting your volumes.

Table 53. Comparing Synchronous Replication and Traditional Replication provides in-depth information about the differences between the two features.

Table 53. Comparing Synchronous Replication and Traditional Replication

Consideration	Traditional Replication	Synchronous Replication (SyncRep)
Typical use case	A point-in-time process that is conducted between two groups, often in geographically diverse locations. Replication provides protection against a regional disaster such as an earthquake or hurricane. Traditional replication has the advantage of providing multiple recovery points. A disadvantage of traditional replication is that the state of the data between recovery points is unknown; if any changes are made to the volume since the last replica was created, they could be lost.	A real-time process that keeps two identical copies of volume data in two different pools within the same PS Series group. Synchronous replication is useful for maintaining two copies of a volume's data in the same data center, or dispersed to two different facilities on the same campus or in the same metropolitan area. An advantage of synchronous replication is that it captures a duplicate copy of every write. One disadvantage is that if an application writes bad data to the volume, the bad data is simultaneously written to both the SyncActive and SyncAlternate volumes.
Recovery time	If a disaster occurs in the primary group, you can promote the replica set on the secondary group to a recovery volume. After the promotion, you must reconfigure initiators to discover and log in to the iSCSI target now hosted by the secondary group, or switch to an alternate set of server resources that have been preconfigured to use the secondary group storage. See "Impact on applications" for more information.	If a disaster involving the active pool occurs, you can manually switch the volume to the alternate pool. After the switch, the alternate pool becomes the active pool and hosts the volume. Host access to the volume is disrupted by the switch, but iSCSI initiators do not need to be reconfigured.
Recovery point	The recovery volume contains point-in-time data that is current as of the most recent replica. Replication can be scheduled to take place as frequently as once every 5 minutes. You can also restore data to the point in time when any previous replicas were created, provided that the replicas have been retained.	Synchronous replication provides a single recovery point: the most recent acknowledged write to the volume.
Network requirements	Replication requires that the network connection between the primary and secondary groups must be able to handle the load of the data transfer and complete the replication in a timely manner.	Because writes are not acknowledged until they are written to both the active and alternate pools, synchronous replication is sensitive to network latency. The network must be able to handle the load of the data transfer from the active pool to the alternate pool and complete the replication in a timely manner, or application performance could suffer.
Snapshots	Replication is functionally similar to snapshots, creating point-in-time copies of the volume. If the `keep failback` option is enabled, the group creates a "failback snapshot" on the primary group every time a replica is created.	Synchronous replication creates snapshots of the volume whenever the SyncActive and SyncAlternate volumes are switched. In addition, you can schedule the creation of snapshots, or create them on demand, as you would with any other volume.

Consideration	Traditional Replication	Synchronous Replication (SyncRep)
		See About Synchronous Replication and Snapshots for more information.
Scheduling	Replication operations can be scheduled using the same mechanism used for scheduling snapshots.	Replication between the SyncActive and SyncAlternate volumes is continuous. Therefore, scheduling synchronous replication is not required.
Pool space requirements	The primary group must have enough space for the volume reserve and local replication reserve, in addition to any snapshot reserve. The secondary group must have enough free space delegated to the primary group for the volume reserve and the replicas that record changes to the volume's data over time.	Both the active pool and the alternate pool must have enough space for the volume and snapshot reserve.
Impact on applications	iSCSI initiators must be reconfigured to connect to the secondary group after the failover, or an alternate set of host resources must be brought online, both of which might cause application disruptions. If you are using the Host Integration Tools, you can coordinate replication with host software to quiesce applications on a schedule and create application-consistent Smart Copies. Replication can help protect against the corruption of application data: depending on when the replica occurred and what your replica retention policies are, you might be able to restore the volume to a point in time before the corruption occurred.	Pool switches might cause disruptions in host access to the volume, but no change to the iSCSI initiator configuration is required to restore access. Writes must be committed to both pools before they are acknowledged to the host, so the application must be able to tolerate whatever additional delay is caused by the simultaneous writes. When synchronous replication is first enabled, or at any other time when data is being written to both copies of the volume to become in sync, performance degradation might occur. This effect is diminished after the volume becomes in sync.
PS Series group requirements	Two PS Series groups, each of which must contain at least one member. The groups' pool configuration is not a consideration.	One PS Series group containing two storage pools, each of which must contain at least one member.

How Synchronous Replication Protects Volume Availability in Different Scenarios

The following scenarios describe how synchronous replication (SyncRep) protects volume availability:

- Normal synchronous replication operation
- SyncAlternate volume unavailable
- SyncActive volume unavailable

Normal Synchronous Replication Operation

In normal synchronous replication operation, in which the volume is in sync, the pools containing the SyncActive and SyncAlternate volumes contain identical data. Volume writes are accepted as shown in Figure 32. Synchronous Replication.

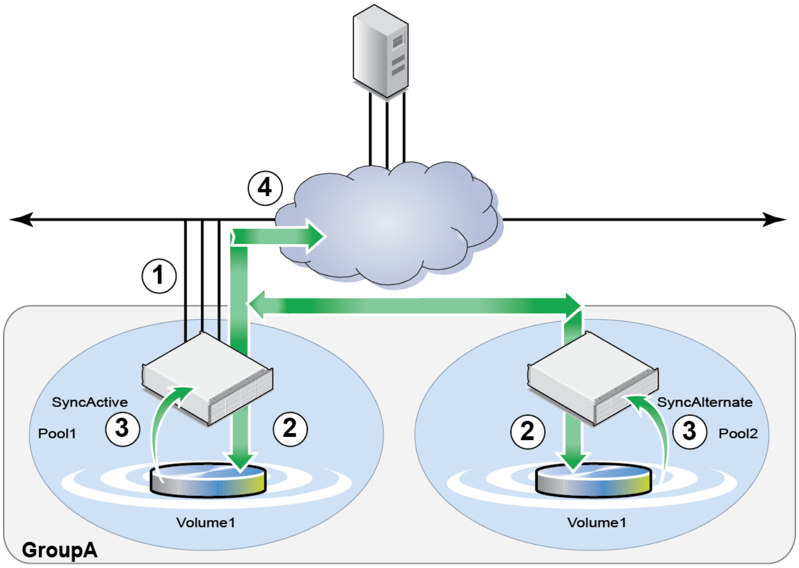

Figure 32. Synchronous Replication

1. The iSCSI initiator sends a write to the group.
2. Data is simultaneously written to both the SyncActive and SyncAlternate volumes.
3. The SyncActive and SyncAlternate volumes confirm to the group that the writes are complete.
4. The write is confirmed to the iSCSI initiator.

SyncAlternate Volume Unavailable

If the group can write to the SyncActive volume, but can no longer write to the SyncAlternate volume, initiator access to the volume continues without disruption, as shown in Figure 33. SyncAlternate Volume Unavailable.

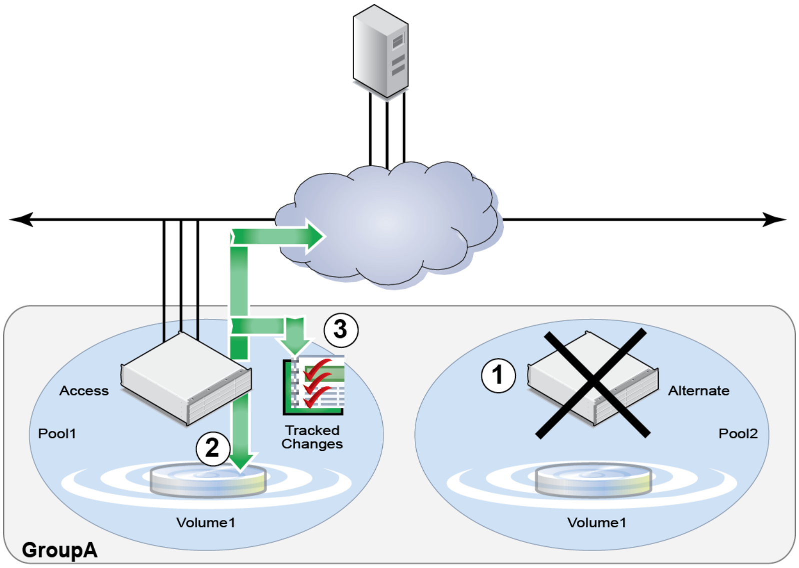

Figure 33. SyncAlternate Volume Unavailable

1. The SyncAlternate volume becomes unavailable.
2. Initiator access to the volume continues without interruption. The volume is out of sync.
3. The group tracks all changes written to the volume.
4. When the SyncAlternate volume becomes available, tracked changes are written to the SyncAlternate volume, as shown in Figure 34. Tracked Changes Written to SyncAlternate Volume.

- Until all tracked changes are written, the data in the SyncAlternate volume is valid only up to the point in time when the volume went out of sync.
- While changes are being tracked or when tracked changes are being written back to the SyncAlternate volume, performance might be temporarily degraded.

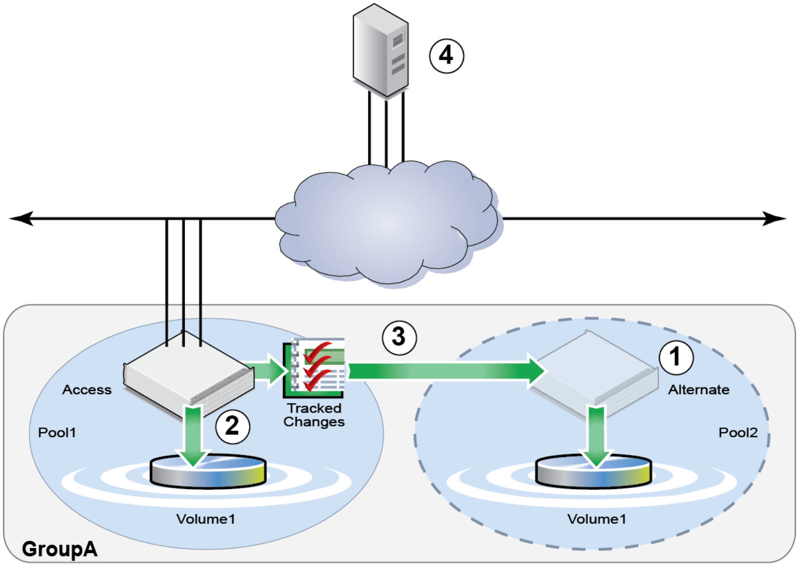

Figure 34. Tracked Changes Written to SyncAlternate Volume

5. When all tracked changes are written, the volume goes back in sync.
6. New writes are simultaneously written to both the SyncActive and SyncAlternate volumes, and normal synchronous replication operations resume.

SyncActive Volume Unavailable

If a malfunction occurs in the SyncActive pool, or some other event has occurred causing the volume to go offline, you can safely switch or fail over to the SyncAlternate volume. See the following sections.

Volume In Sync

If the volume is in sync, you can switch to the SyncAlternate volume. Although host access to the volume is disrupted during the switch, no initiator changes are required.

Volume Out of Sync

1. Log in to the Group Manager GUI using an IP address that belongs to a group member in the pool containing the SyncActive volume. Do not use the group IP address.
2. Set the volume offline.
3. Verify that no initiators are attempting to reconnect to the volume.
4. Attempt to fail over the volume.

   CAUTION: Failovers should be performed only under extraordinary circumstances, or in cases where you can be sure that the failover will not destroy any data.

   If you fail over to the SyncAlternate while the volume is out of sync, any changes written to the volume since it went out of sync will be written to a snapshot. As with other snapshots, this snapshot can be cloned or restored, but is also subject to deletion by the group's snapshot retention policy. If the snapshot is deleted, its data will be lost and cannot be recovered.

   If the failover was successful:

   - Log in to the Group Manager GUI using an IP address that belongs to a group member in the pool containing the SyncAlternate volume. Do not use the group IP address.

   If the failover was not successful:

- Disconnect the SyncActive volume, as documented in the Group Manager online help. If the disconnect fails, try logging in to a different member in the SyncActive pool.

  NOTE: If the disconnect operation will not succeed from any member in the SyncActive pool, contact your Dell support provider for assistance.

- Log in to the Group Manager GUI using an IP address that belongs to a group member in the pool containing the SyncAlternate volume. Do not use the group IP address.
- Fail over the volume.

5. Set the volume back online, restoring initiator access. The new SyncActive volume (formerly SyncAlternate) accepts writes and tracks changes made while the original SyncActive volume is offline. Depending on host configuration and application requirements, you might have to restart hosts that were connected to the volume.
6. After access to the volume has been restored, the system performs the following steps:

   a. When the original SyncActive volume is available again, it becomes the SyncAlternate volume.

   b. Synchronous replication resumes. The volume is out of sync, and tracked changes are written to the new SyncAlternate volume.

   c. When the volume becomes in sync again, new writes are simultaneously written to both pools. You can switch back to the original pool configuration, or continue using the new pool configuration.

Requirements for Using Synchronous Replication

Before you can configure a volume to use synchronous replication, verify that the following requirements are met:

- Two pools, each containing at least one member
- Adequate network bandwidth between pools
- Free space in each pool to accommodate the volume and snapshot reserve for the volume

Also note the following restrictions:

- You cannot enable synchronous replication on a volume for which traditional replication is configured, and you cannot enable traditional replication on a volume for which synchronous replication is configured.
- You cannot enable synchronous replication on a volume that is bound to a group member. See About Overriding Automatic Load Balancing for more information about binding and unbinding volumes.
- To support up to 32 volumes for synchronous replication, a PS Series group must contain only PS6cxx0 arrays. If the group contains one or more PS4xx0 arrays, only 4 volumes are supported.

Synchronous Replication States

Synchronous replication states for a volume include:

- in sync

  SyncActive and SyncAlternate copies of a volume contain the same data.
- paused

  Administrator has paused synchronous replication. While synchronous replication is paused, the volume is still online, and initiators can connect to and write to the SyncActive volume.

  An administrator might pause and later resume synchronous replication. For example, this situation could happen in a maintenance window during which the SyncAlternate volume is taken offline. If data is written to the volume while synchronous replication is paused, it is written only to the SyncActive volume, and the two volumes are out of sync. The group tracks all volume writes while synchronous replication is paused and, when the administrator resumes synchronous replication, writes the tracked changes to the SyncAlternate volume.
- out of sync

  SyncActive volume and SyncAlternate volume might not contain the same data; the SyncActive volume contains the most recent data.

A volume can become out of sync if synchronous replication is paused, or if one of the volumes becomes unavailable or has no free space. The volume can become out of sync when the snapshot reserve in the SyncAlternate volume is full, but only when the snapshot space recovery policy sets volumes offline when the snapshot reserve is depleted.

Whenever the volume's state changes to paused or out of sync, the group creates a snapshot of the volume that reflects the volume's contents at the point in time when the state changed. This snapshot resides in the active pool. If not enough space in the snapshot reserve is available for the snapshot, the group does not create the snapshot.

NOTE: When synchronous replication is first enabled, or at any other time when data is being written to both volumes to become in sync, performance degradation might occur. This effect increases with the quantity of tracked changes, but it is significantly reduced after the volumes become in sync. Also, depending on the quantity of tracked changes, activity within the group, and available network bandwidth, an extended period of time might elapse before the two volumes become in sync again. The Group Manager GUI displays the status of this operation.

About System Snapshots and SyncRep

System snapshots are a special type of snapshot that are automatically created by the system during management of synchronous replication (SyncRep) volumes. System snapshots reside in the same snapshot reserve space as other snapshots. Although you cannot modify their attributes, they are visible in the Group Manager GUI or CLI, and you can create clones of them. You can clone a system snapshot if you need access to the volume's data as it existed at the point in time where the volume went out of sync.

System snapshots are of two types, protected and failback.

Protected System Snapshots

Protected snapshots are created before the system writes tracked changes to the SyncAlternate volume to keep the volume in sync. Protected snapshots reside in the alternate pool. They represent the state of the SyncAlternate volume at the point in time when the volume went out of sync.

You cannot delete protected snapshots; the system deletes them automatically after the volume becomes in sync.

If not enough space in the snapshot reserve is available for the protected snapshot, the system will not be able to put the volume back in sync unless the group's snapshot space recovery policy has snapshot borrowing enabled.

Failback System Snapshots

Failback snapshots are created in the active pool when the volume is paused or goes out of sync. If not enough snapshot space is available to contain the failback snapshot, you cannot fail back to the SyncActive volume unless the group's snapshot space recovery policy has snapshot borrowing enabled.

About Synchronous Replication and Snapshots

You can create snapshots of synchronous replication volumes. You can access the data in the snapshots by setting the snapshot online or cloning it.

With synchronous replication volumes, the snapshot reserve percentage you specify for a volume is applied to both the SyncActive and SyncAlternate pools. For example, if you are creating a 20GB volume and setting the snapshot reserve to 200 percent, you need to have:

- 20GB for the volume reserve in the SyncActive pool
- 20GB for the snapshot reserve in the SyncActive pool
- 20GB for the volume reserve in the SyncAlternate pool
- 20GB for the snapshot reserve in the SyncAlternate pool

NOTE: If data reduction has been enabled on the volume, snapshot reserve is permanently disabled.

When you create a snapshot of a volume for which synchronous replication (SyncRep) is enabled, the snapshot resides in the pool that is the SyncActive pool at the time the snapshot is created. If the SyncActive pool is switched, existing snapshots remain in the pool in which they were created, and subsequent snapshots reside in the new SyncActive pool.

You can perform the same operations on any of the volume's snapshots, regardless of whether they reside in the SyncAlternate or SyncActive pool, with one exception: Only snapshots residing in the SyncActive pool can be restored. To restore a snapshot in the SyncAlternate pool, you must first switch pools.

For example, assume a synchronous replication volume named `Volume1` has a SyncActive pool named `PoolA` and a SyncAlternate pool named `PoolB`. The group administrator takes 5 snapshots of the volume. All of the snapshots reside in `PoolA`. Later, the group administrator switches pools, and creates 3 snapshots. These new snapshots reside in `PoolB`. Now, 5 snapshots are in `PoolA` and 3 snapshots are in `PoolB`. The group administrator subsequently switches back to the original configuration and creates 2 additional snapshots that reside in `PoolA`. Now, 7 snapshots are in `PoolA` and 3 snapshots are in `PoolB`.

Automatic Snapshot Creation

The group creates snapshots of synchronous replication volumes in the following circumstances:

- Anytime an administrator performs a pool switch, failover, or disconnect
- Whenever the volume's state changes to paused or out of sync
- Just before the group begins writing tracked changes to a the SyncAlternate pool in an out-of-sync volume

These snapshots are called protected system snapshots and failback system snapshots, and they are visible in the GUI and CLI. System snapshots reflect the volume's data at the point in time when the operation was performed or the volume's status changed. Depending on their type, system snapshots can reside in either the SyncActive or the SyncAlternate pool. If the SyncActive pool snapshot reserve does not have enough space, the group does not create a snapshot. Not creating a snapshot does not have any impact on the operation or status change, or on the volume's availability.

Snapshots created during a pool switch are deleted after the switch has completed successfully and the volume's status returns to in sync.

If the SyncAlternate pool snapshot reserve does not have enough space for a system snapshot, the group cannot put an out-of-sync volume back in sync. To ensure that a volume can be put back in sync, Dell recommends that snapshot borrowing be enabled in groups where synchronous replication is in use, and that the snapshot reserve for synchronous replication volumes be set to 100 percent or more of the volume's size.

About Synchronous Replication Switches and Failovers

You can change iSCSI initiator access from the SyncActive volume to the SyncAlternate volume using either switches or failovers.

NOTE: Although no changes to the volume's iSCSI initiator configuration are required, the change might disrupt host connectivity to the volume. In some cases, you might have to restart host applications, iSCSI initiators, or operating systems for hosts to reconnect to the volume. Consequently, Dell recommends that you perform synchronous replication (SyncRep) switches and failovers during periods of downtime or minimal activity.

SyncRep Switches

A pool switch for a synchronous replication volume swaps the roles of the SyncActive and SyncAlternate volumes. As part of the operation, the SyncActive volume becomes the SyncAlternate volume and the SyncAlternate volume becomes the SyncActive volume. When a pool switch takes place, iSCSI initiators are temporarily logged out of the volume. They can log back in again after the switch is complete. Depending on the initiator configuration, the initiators might automatically attempt to reconnect. After the failover or switch completes, the volume is automatically set online and synchronous replication resumes.

You can switch the synchronous replication configuration if a failure is imminent for the active pool, or if maintenance needs to be performed on the array hardware in the active pool. You can also switch pools at any time, even if a failure has occurred, provided that the volume is in sync. Aside from the brief period when the volume is offline during the switch, switching eliminates downtime during a maintenance window on the active pool.

SyncRep Failovers

If the volume is out of sync, the option to switch the pools is replaced with the option to fail over to the alternate pool. You can fail over to the alternate pool only when the volume is offline. As with a pool switch, a pool failover pauses synchronous replication. When the failover is complete, you must manually set the volume online and resume synchronous replication.

CAUTION: Failovers should be performed only under extraordinary circumstances, or in cases where you can be sure that the failover will not destroy any data.

If you fail over to the SyncAlternate volume while the volume is out of sync, any changes written to the volume since it went out of sync will be written to a snapshot. As with other snapshots, this snapshot can be cloned or restored, but is also subject to deletion by the group's snapshot retention policy. If the snapshot is deleted, its data will be lost and cannot be recovered.

When the condition necessitating the switch or failover has been resolved, the administrator can either switch back to the original configuration (if the volume is in sync) or leave the pools in the new configuration.

About Synchronous Replication Volume Collections

Synchronous replication (SyncRep) volume collections are a unique type of volume collection that can only contain synchronous replication-enabled volumes. As with other volume collections, you can clone or take snapshots of synchronous replication collections. However, synchronous replication collections have the additional capability of supporting synchronous replication operations.

In synchronous replication volume collections, all volumes in the collection behave as a group. The volumes simultaneously reside in the same active and alternate pools. When synchronous replication switches are performed on the collection, they take place for all volumes in the collection.

You might want to use synchronous replication collections in cases where you have a group of volumes that all must contain data that is current up to the same point in time, such as with an application or host system that is using multiple volumes. If a single volume in the collection goes out of sync, the group takes snapshots of all volumes in the collection, creating a common point in time to which all volumes in the collection can be restored.

The following requirements and restrictions apply when working with synchronous replication volume collections:

- Synchronous replication collections are considered to be in sync only when all volumes are in sync, and go out of sync if a single volume within the collection ceases to be in sync.
- Synchronous replication switches cannot be performed on individual volumes in the collection.
- If you create a collection composed of synchronous replication volumes, synchronous replication is not automatically enabled for the collection. You must still enable synchronous replication for the collection.
- All volumes in a synchronous replication collection must have synchronous replication enabled before being added to the collection.
- Synchronous replication-enabled volumes can be in multiple collections, but they can belong to only one synchronous replication collection at a time.
- A volume must be in sync before it can be added to a synchronous replication collection.
- You cannot change the pool assignments of volumes in synchronous replication collections.
- You cannot assign a volume to a synchronous replication collection while you are changing its pool assignment.
- All volumes in a synchronous replication volume collection must have the same active and alternate pools. You cannot add volumes to the collection if they have pool assignments that differ from the assignment specified for the volumes in the collection. To add these volumes to the collection, you must first change their pool assignments so that they match that of the collection.

For example, assume you have a collection for which the active pool is Pool-A and the alternate pool is Pool-B, and Pool-C is a third pool in the group. Any synchronous replication-enabled volumes that use Pool-C must have their pool assignments changed to be using Pool-A and Pool-B before they can be added to the collection.

- When you create the collection, active and alternate pools for the collection's volumes are chosen based on the pool assignment of the first volume added to the collection. If you create the collection with the GUI, the group switches the active and alternate pools as necessary so that all volumes in the collection have the same pool assignments. In the CLI, you must manually set the pool assignments before creating the collection.

  For example, assume you have a collection for which the active pool is Pool-A and the alternate pool is Pool-B. Any volumes that have Pool-B for their active pool and Pool-A for their alternate pool will be switched when they are added to the collection.

  NOTE: Pool switches briefly take the volume offline, which might disrupt iSCSI connections to the volume. Consequently, Dell recommends that synchronous replication volume collections should be created only during periods of downtime or low activity.

- You can pause synchronous replication for one volume in the collection or for the entire collection. If you pause synchronous replication for a single volume in the collection, the sychronous replication status of all other volumes in the collection is also shown as out of sync, even though data is still written to their active and alternate pools. The system brings the collection back in sync after synchronous replication is resumed for all volumes.

About Using Thin Clones and Templates with Synchronous Replication

If you are going to use synchronous replication with thin clones and template volumes, the following restrictions and considerations apply.

Thin Clones

- When you create a thin clone of a template volume for which synchronous replication is enabled, the thin clone resides in the template volume's active pool.
- You must manually enable synchronous replication for thin clones. Synchronous replication is not automatically enabled for the thin clones of a synchronous replication-enabled template.
- You can enable synchronous replication for a thin clone if the following conditions are met:
  - Synchronous replication is enabled for the thin clone's template.
  - The initial copy of the template's data from the active pool to the alternate pool has completed.
- All thin clones of a template volume must reside in the same pools as the template. However, their active and alternate pool assignments might differ.

  For example, a template volume might have `PoolA` for its active pool and `PoolB` for its alternate pool. Any or all of its thin clones might have `PoolB` for its active pool and `PoolA` for its alternate pool.
- Performing a switch on a thin clone neither switches the pools of any other thin clones that share the same template, nor switches the template's pools.
- Disabling synchronous replication for a thin clone does not change the synchronous replication configuration of its template, or that of any other thin clones that share the template.

Templates

- You cannot switch the pool of a template until the initial copy of the template's data from the active pool to the alternate pool has completed.
- If synchronous replication is enabled for a template volume, you can detach one of its thin clones at any time, provided that both the active and alternate pools are available. If they are not available, you must disable synchronous replication for the thin clone before detaching it.
- You can convert a synchronous replication template back into a volume only when the following conditions are met:
  - The initial data copy from the SyncActive volume to the SyncAlternate volume has completed.
  - The active and alternate pools are available. If they are not both available, you must disable synchronous replication for the template before performing the conversion.
- Before disabling synchronous replication for a template, you must first disable synchronous replication for each of its thin clones.
- If you disable synchronous replication on a template, a copy of the template remains in the alternate pool.

- Before converting a synchronous replication volume into a template, it must be in sync. To convert an out-of-sync volume to a template, you must first disable synchronous replication for the volume.

Configure Synchronous Replication (SyncRep) on a Volume

Before you can configure synchronous replication for a group, the group must include at least two different storage pools.

See the following sections for more information: Create an Empty Storage Pool and Create a Storage Pool from an Existing Member.

NOTE: You cannot enable synchronous replication on a volume for which traditional replication is configured, and you cannot enable traditional replication on a volume for which synchronous replication is configured. For instructions on disabling replication for a volume, see Disable Replication.

To configure synchronous replication (SyncRep) on a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Configure SyncRep** to open the Configure SyncRep wizard.
4. In the **General properties** section, specify the alternate pool for the volume.

   The volume's pool is automatically selected for its active pool.
5. Click **Next** to go to the Summary step of the wizard.
6. Verify that the information contained in the summary is correct.
7. Click **Finish** to complete the wizard and enable synchronous replication.

NOTE: After you have enabled synchronous replication, it might take several minutes for the volume's active and alternate pools to become in sync.

Disable Synchronous Replication (SyncRep) for a Volume

NOTE: You cannot disable synchronous replication if the active pool is not available.

To disable synchronous replication for a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Disable SyncRep**.

   When prompted, confirm the action.

CAUTION: If you disable synchronous replication for a volume, the volume will reside in the active pool. All of the volume's data in the alternate pool will be deleted. However, if you want the volume to reside in the alternate pool, switch to the alternate pool before disabling SyncRep.

Monitor Synchronous Replication (SyncRep) Volumes

To monitor all volumes configured for synchronous replication (SyncRep):

1. Click **Monitoring** in the navigation menu.
2. In the SyncRep section of the navigation tree, select **SyncRep Volumes**.

The SyncRep panel displays a list of volumes.

Pause Synchronous Replication (SyncRep)

To pause synchronous replication (SyncRep) for a volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Pause SyncRep**.

   When prompted, confirm the action.

While synchronous replication is paused:

- The volume's SyncRep status in the General Volume Information panel indicates that it is paused.
- If any data is written to the volume while synchronous replication is paused, the writes are tracked. The tracked changes are written to the SyncAlternate volume when synchronous replication is resumed.
- If you pause synchronous replication for a single volume in a synchronous replication collection, the synchronous replication status of the paused volume is paused, and the other volumes in the collection are shown as being out of sync, even though data is still written to their active and alternate pools. The system brings the collection back in sync after synchronous replication is resumed for all volumes.

Resume Synchronous Replication (SyncRep)

To resume synchronous replication (SyncRep) after it has been paused:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Resume SyncRep**.

After synchronous replication has been resumed:

- The group copies to the SyncAlternate volume any changes that were made while synchronous replication was paused.
- The volume's SyncRep status in the General Volume Information panel indicates that the volume is out of sync until all outstanding writes have been committed to the SyncAlternate volume. When all writes have been committed, the status changes to in sync.

Enable Synchronous Replication (SyncRep) for a Volume Collection

To enable synchronous replication (SyncRep) for a volume collection, all volumes in the collection must:

- Have synchronous replication enabled
- Be in sync
- Use the same two storage pools
- Not belong to another SyncRep collection
- Not be moving from one pool to another

To enable SyncRep for a volume collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection from the list.
3. In the Activities panel, click **Enable SyncRep for collection**.

Disable Synchronous Replication (SyncRep) for a Volume Collection

To disable synchronous replication (SyncRep) for a collection:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select a collection from the list
3. In the Activities panel, click **Disable SyncRep for collection**.

Change the Pool Assignment of a Synchronous Replication (SyncRep) Volume

To change the pool assignment for a synchronous replication volume, use the following steps.

NOTE: The same free space requirements apply to changing the pool containing the SyncActive or SyncAlternate volume that also apply when moving a volume for which synchronous replication is not enabled.

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Move volume** to open the Move Volume dialog box.
4. In the **Volume to Move** section of the dialog box, select the SyncActive or SyncAlternate volume.
5. In the **Destination Pool** section of the dialog box, specify the new pool for the selected volume.

   NOTE: The destination pool cannot be either of the pools that are being used for either the SyncActive or SyncAlternate volume.
6. Click **OK** to change the pool assignment. Changing the pool assignment does not disrupt iSCSI initiator access to the synchronous replication-enabled volume.

View the Distribution of a Volume Across Pools

To view the distribution of a synchronous replication (SyncRep) volume across its active and alternate pools:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. The **Volume and Snapshot Space** panel displays the volume's distribution across its pools.

About Switching and Failing Over SyncRep Pools

During failovers and switches, the volume is temporarily set offline and synchronous replication (SyncRep) is paused. After the failover or switch completes, the volume is automatically set online and synchronous replication resumes.

Although no changes to the volume's iSCSI initiator configuration are required, the change might disrupt host connectivity to the volume. In some cases, you might have to restart host applications, iSCSI initiators, or operating systems for hosts to reconnect to the volume. Consequently, Dell recommends that synchronous replication switches and failovers be performed during periods of downtime or minimal activity.

Switch from SyncActive to SyncAlternate

To switch a volume from the active pool to the alternate pool:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Switch to SyncAlternate**.

   You will be prompted for confirmation of the action. After the switch occurs, the Volume and Snapshot Space panel reflects the pool change.

Fail Over from SyncActive to SyncAlternate

The steps for failing over a volume are similar to switching a volume; when the volume is out of sync, the option to switch the volume is replaced with an option to fail over the volume. Failovers can be performed only if the volume is offline.

CAUTION: Failovers should be performed only under extraordinary circumstances, or in cases where you can be sure that the failover will not destroy any data.

If you fail over to the alternate pool while the volume is out of sync, any changes written to the volume since it went out of sync will be written to a snapshot. As with other snapshots, it can be cloned or restored, but is also subject to deletion by the group's snapshot retention policy. If the snapshot is deleted, its data will be lost and cannot be recovered.

If you fail over to the alternate pool while the active pool is disconnected, all unreplicated data will be lost and unrecoverable.

To fail over a volume from the active pool to the alternate pool:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the Activities panel, click **Failover to SyncAlternate**.

   You will be prompted for confirmation of the action. After the failover occurs, the Volume and Snapshot Space panel on the Status tab reflects the pool change.

Switch Synchronous Replication (SyncRep) Collection Pools

To switch a synchronous replication (SyncRep) collection to the SyncAlternate pool:

1. Click **Volumes**.
2. Expand **Volume Collections** and then select the collection from the list.
3. In the Activities panel, click **Switch collection to SyncAlternate**.

NOTE: If the collection is out of sync, the Switch option is replaced with an option for failing over the collection.

Disconnect the SyncActive Volume

If the pool containing the SyncActive volume becomes unavailable to initiators, or communications between the SyncActive and SyncAlternate volumes are disrupted, you can restore volume connectivity by disconnecting the SyncActive volume and failing over to the SyncAlternate volume.

NOTE: Disconnecting the SyncActive volume takes the volume offline. The operation cannot be reversed, and volume data is unavailable until you switch or fail over to the SyncAlternate volume.

To disconnect the SyncActive volume:

1. Click **Volumes**.
2. Expand **Volumes** and then select the volume.
3. In the SyncRep section of the Activities panel, click **Disconnect SyncActive**.

   The volume is taken offline.

19

About Self-Encrypting Drives (SEDs) and AutoSED

A self-encrypting drive (SED) performs Advanced Encryption Standard (AES) encryption on all data stored within that drive. SED hardware handles this encryption in real-time with no impact on performance. To protect your data, a SED will immediately lock itself whenever it is removed from the array (or otherwise powers down). If the drive is lost or stolen, its contents are inaccessible without the encryption key.

The Dell AutoSED feature provides the benefits of SED security with no effort on the part of the administrator. You do not need to configure or set up drives, manage encryption, or install a Key Management Service (KMS). Everything is handled by AutoSED.

AutoSED operates at the level of the physical disk drives within an individual member. SEDs cannot be used to encrypt volumes, in the sense of securing each iSCSI volume with its own key. SEDs also cannot provide security across members, so it is up to the administrator to ensure that SED members and non-SED members are properly deployed.

Using AutoSED is effortless, but it is important to understand what protection AutoSED provides and what protection it does not provide.

Central to the AutoSED security model is the concept of a SEDset. Similar to how RAID groups drives into a RAIDset for redundancy, AutoSED groups drives into a SEDset for security. Each member of a group has one SEDset that spans all active drives in the member. The SEDset cannot be unlocked unless it is sufficiently intact, which means that at least half of its drives are present.

Scenarios Covered by AutoSED

- Loss of a drive — When a drive leaves the SEDset (whether by failure, removal, or otherwise), the drive immediately locks itself. Its contents are inaccessible without the encryption key, which is owned by the SEDset. At the same time, the SEDset immediately resecures itself to exclude the departed drive, preventing access to the key.
- Loss of fewer than half the drives — When fewer than half the drives in the SEDset are removed, the SEDset remains intact and resecures itself to exclude all the removed drives. The removed drives are locked, and have no access to the SEDset key.
- Loss of other array components — The SEDset key resides wholly within the drives. The key cannot be found in the flash cards, channel cards, midplane, chassis, or any other component, including the controllers and controller memory.

Scenarios Not Covered by AutoSED

- Loss of the entire array — A SEDset is a self-contained apparatus, which is why the array can unlock itself with no external assistance. A stolen array will continue to unlock itself, just as it did before it was stolen.
- Loss of half the drives — Security might be compromised if half (or more) of the drives are removed at one time. These drives can be combined into an intact SEDset of their own, which will automatically unlock itself.
- Insider attack — Any person who possesses the administrator password can access any volume on the array, or change ACLs to allow others to do the same. Similarly, a compromised host can access volumes that the host is authorized to access. SED is irrelevant in these cases.
- Data in flight — SEDs provide no protection for data in flight on the network. IPsec should be used to provide secure connections to the array.
- Tampering with array hardware — AutoSED is not resistant to modified firmware, hardware probes and other snooping devices, or the removal of a drive without loss of power to that drive.

About Self-Encrypting Drives (SED)

SEDs (self-encrypting drives) are disk drives that use an encryption key to secure the data stored on the disk. This encryption protects the PS series array from data theft when a drive is removed from the array.

SED operates across all disks in an array at once. If one drive in a RAID set is removed from the array, a new set of encryption key shares is generated automatically and shared among the remaining disks. If a second drive is removed from the same RAID set, another set of encryption key shares is generated.

SED drives are configured at the factory. When the drives are installed into an array, the array automatically detects the new SED drives and locks them. This process is automatic; the GUI has no user controls for SED.

All of the drives in an array, including spares, must be of the same type and model, and must be running PS Series firmware 6.0 or higher. A SED drive installed into a mixed-disk configuration, or a configuration containing unencrypted drives, operates as an unencrypted disk. Likewise, a pool consisting of all SED drives might replicate to a pool with only a few SED drives or no SED drives at all.

NOTE: SED drives are identified in the GUI with a gold key icon.

How Key Shares Work

Each array has an overall shared encryption key that protects data on all of the disks in that array.

The shared encryption key is not stored in any one location on the array. Instead, the key is divided into portions called *key shares*. The number of key shares generated corresponds to the number of drives in the array (except for spares or other drives not used by the array). The key shares are distributed across all non-spare disks used in the RAID configuration. If your array has *n* non-spare disks, you must have (*n*+1)/2 of the key shares to unlock the data on the disks. If you are missing one or more of the key shares, you will not be able to recover the data.

You can back up the disk encryption key shares. Key shares are backed up in groups of three files. To unlock the array, you need to supply two backup shares. Under normal operation, the keys are not necessary because the data is redundant; however; they might be useful in the event that a disk needs to be sent to a data recovery service. Use the Maintenance tab to back up the disk encryption key shares.

How Self-Encryption Protects Data

To understand how SED protects your data, you should understand the types of threats to data that SED cannot protect against.

Each individual drive has its own secure PIN, which is local to that drive and not shared. If a drive fails, this PIN is needed to unlock and recover the information on that drive. All drives also have the key shares. During normal operation, the data redundancy across the array prevents individual drives from becoming single points of failure.

SED protects against data theft in the following circumstances:

- Loss or physical removal of fewer than half of the drives installed in an array (not counting the spares). This number includes drives that are removed for reuse elsewhere, as well as drives lost due to theft. The SED key remains secure and the disk encryption remains unbroken as long as more than half of the drives remain in the array.
- “Cold memory” attacks that attempt to extract data from system DRAM by powering down the array or removing the drive. SED immediately erases any data in DRAM when power is lost.
- Because SED drives are configured and shipped from the factory, SED protects the entire drive immediately. You cannot configure individual volumes for self-encryption.

CAUTION: You must have (*n*+1)/2 of the current key shares to unlock the array, if the array has *n* drives installed. If you lose the keys, then the data on the drive will be irrevocably lost.

SED cannot protect against the following threats:

- Loss of the entire array, or simultaneous loss of half of the drives in the array. If half of the drives on an array are lost, the data on those drives is compromised. The locking mechanism for the remaining drives is also compromised, leaving the data exposed. If more than half of the drives are lost, the array is rendered inoperable. In the case of a RAID-10 configuration, loss or theft of one drive in each mirrored pair could result in the data on the entire array being exposed and the RAID sets rendered inoperable.
- If the drive is removed from the array but power is maintained to that drive, the drive remains unlocked. Data locking occurs when power to the drive is lost.
- It is possible to recover the encryption key data from the SAS link if it is compromised ("snooped") during the process of unlocking a locked drive.
- Insider attacks. For example, an administrator account can still change the ACLs on volumes, allowing different initiators to access the data on the disks. Administrators can read any volume, grant permission to anyone to read any volume, can obtain the key shares, and so on.

About SED Members in a Group

AutoSED operates at the member level, not at the group level. SED management is completely automatic within a member. Within a group, the administrator is responsible for deploying SED members and non-SED members properly.

A pool must contain only SED members to be secure. A gold key icon indicates that the pool is fully encrypted. Volumes created in this pool are secure.

Though it is permitted to mix SED members and non-SED members in the same pool, it is not recommended. An alert icon indicates that a pool is only partially encrypted.

Mixed pools are intended only to simplify the gradual upgrade of a pool from non-SED to SED. No security benefits should be expected until the entire pool contains only SED members. Creating a volume in a mixed pool provides no security, because some or all of the volume might reside on non-SED members.

Examples

1. A group contains two pools of non-SED members. A new SED member is brought online.

   Add the SED member to its own pool. Volumes that you create in this pool will be secure. Volumes that you create in the non-SED pools continue to be insecure.
2. A group contains a pool of three SED members. A non-SED member is added to this pool.

   Every volume in this pool immediately becomes insecure. Because of Dell EqualLogic's automatic load balancing, volumes and snapshots will be spread across all the members, including the non-SED member. To return the pool to secure operation, delete the non-SED member or move it to a different pool.
3. A group contains a pool of three non-SED members. The intention is to migrate the pool to secure operation.

   Add enough SED members to the pool to hold your data. Delete the non-SED members or move them to a different pool. During this time, the pool is still not secure. When the process completes, the gold key icon signifies that the pool (and all of its volumes) are now secure.

Back Up a Self-Encrypting Drive (SED) Key

The AutoSED machinery remains functional even when severe failures have taken the array offline. The backup is needed only in exceptional circumstances, such as the loss of more than half the drives from an array.

The SED key is never explicitly revealed as part of the backup process. Rather, it is encrypted into a set of three unique backup units. Any two backup units from the same backup set can be combined to decode the key. Although the key never changes (unless the member is reset), each backup set is unique. No two sets are alike, and backup units from different sets cannot be combined.

The array automatically creates and presents a backup set during initial setup, when the RAID policy is configured. Additional backup sets can be manually requested at any time.

During normal operation, the array has the information it needs to operate SED disks. The key shares are stored across the array on the non-spare disks. If a disk fails and is replaced by a spare, the configuration generates a new set of key shares, and the original key shares are discarded.

If a SED disk goes offline due to power failure, removal from the array, or disk failure, the disk is automatically locked, and any data residing in memory about that disk drive is automatically wiped. To recover the data on that disk, you must provide two of the three key shares to unlock the disk. Backing up the key shares ensures that you have current copies in case you need to recover the data on a locked disk.

1. Click **Group** and then expand **Members**.
2. Select the name of the member whose encryption key you want to back up.
3. Click the **Maintenance** tab.
4. In the Disk Encryption panel, click the **Encryption Key Shares** button.
5. Enter the administrative password in the dialog box. The **Information** dialog box lists the names and code string of each key share.
6. To download all three key shares (backup units) as individual text files, click **Save all...** and choose the location where you want to store them. All three file names have the format *membername-keyshare-n*, where*n* stands for 1, 2, or 3.
7. Use the **Copy** buttons above each key share to copy the individual key share (backup unit) and paste it into a file, if desired. Select **Copy all** if you want to copy all three key shares to the clipboard.

NOTE: If you generate a second set of key shares, the first set is not invalidated. Generating a second set of key shares, therefore, does not protect the key shares from being compromised.

Self-Encrypting Drives (SED) Frequently Asked Questions (FAQ)

Why are my backups always different?

Although the encryption key never changes, the backup looks different each time it is generated. The three backup units are cryptographic images of the key, and are never generated the same way twice.

Why is a secure-erase command not available?

The command is not needed. Whenever it is safe to erase a drive, AutoSED will always do so, without intervention. A manual secure-erase is never necessary, so no command is provided to perform it.

NOTE: Secure-erase is also known as cryptographical erase or crypto-erase.

What is the difference between a locked drive and a securely erased drive?

Data that is locked is inaccessible without the SEDset key. Data that is securely erased has been cryptographically destroyed.

I accidentally reset an SED array. What can I do?

Nothing. Every drive in the member has been securely erased, and the data has been cryptographically destroyed. Recovery is impossible.

What if the entire array is stolen?

Security is compromised. The array will unlock itself when it boots, as it did before it was stolen.

What if the grpadmin password is stolen?

Security is compromised. The adversary can connect to the array over the network and read the data. SED is irrelevant in this case.

Is it safe to discard or return a locked SED?

Yes. Any data that you have written to the drive will be locked and inaccessible. When you return a drive to Dell, the only information that remains readable are its operating statistics (S.M.A.R.T. data), its RAID type, and its hardware error logs.

Can I add SEDs to a non-SED array, or vice versa?

No. Do not mix SEDs and non-SEDs in the same array. If mixed drives are detected while the array is booting, the array will halt until the incorrect drives are removed. If mixed drives are detected while the array is operating, the incorrect drives will be shown as unauthorized.

Does a SED system also use RAID?

Yes. Each drive in a SED-equipped array is managed by both AutoSED and RAID. The SEDset governs the locking of data, and the RAIDset governs the data itself.

Does SED encrypt my volumes?

No. SEDs cannot be used to encrypt volumes, in the sense of securing each iSCSI volume with its own key. AutoSED operates at the level of the physical disk drives within an individual member.

If I create a new set of backup units, does the new set invalidate the previous set of backup units?

No. Generating a new set of backup units does not affect previously created backup sets.

Self-Encrypting Drives (SED) Examples

1. SED array is operating normally. Then, a drive is removed (or fails).

   Security is not compromised. The drive immediately locks itself. Its contents are inaccessible without the SEDset key. The SEDset also resecures itself to exclude the removed drive. Therefore, the drive can be safely repurposed, discarded, or returned to Dell.
2. SED array is operating normally. First, drive #1 is removed (or fails). Later, drive #2 is removed (or fails).

   Security is not compromised. When drive #1 is removed, the drive locks itself, and the SEDset is resecured to exclude drive #1. When drive #2 is removed, it also locks itself, and the SEDset is resecured again to exclude drive #2. As a result, both drives now exclude each other. Both drives can be safely repurposed, discarded, or returned to Dell.
3. SED array is operating normally. Then, half of the drives are removed one at a time.

   Security is not compromised. As in Example 2, each drive immediately locks itself upon removal, and the SEDset resecures itself each time. As a result, all the removed drives exclude each other, and cannot be used to construct an intact SEDset despite having enough drives to do so.
4. SED array is operating normally. Then, half of the drives are removed at the same time.

   Security is compromised. Although the removed drives immediately lock themselves, the adversary now possesses enough drives to construct an intact SEDset, which will unlock itself when booted in an appropriate array.

   Removal of half of the drives always causes the array to stop functioning and alert the administrator, with one exception. A RAID-10 configuration remains operational if one drive from each mirror pair remains in the array.
5. SED array is powered off. Then, half the drives are removed one at a time.

   Security is compromised. Because the array is not operating, the SEDset cannot resecure itself. The adversary now possesses enough drives to construct an intact SEDset.
6. Two SED arrays, X and Y, are operating normally. A drive is moved from array X to array Y.

Security is not compromised. Array Y cannot unlock the drive because it needs the SEDset key from array X. The drive can be manually converted to a spare, and doing so will instantly erase it.

7. SED array is operating normally. A drive and a controller are removed.

   Security is not compromised on the drive. The SEDset key cannot be found on the controller, even if it is pulled from a running system. However, cached data might be found in the controller's battery-backed RAM, which is not protected by SED or any other encryption.

8. SED array with 16 slots is populated with 8 SEDs. Then, 8 new SEDs are added.

   Assume the array includes 6 active drives and 2 spares. Initially, the SEDset spans the 6 active drives, so 3 drives must be lost before the key is compromised. As new drives are inserted, the SEDset resecures itself with each addition. Eventually, the SEDset spans all 14 active drives, so 7 drives must be lost before the key is compromised.

Self-Encrypting Drives (SED) Advanced Encryption

Advanced encryption for SEDs includes the following methods:

- Media encryption key and access key
- Threshold secret sharing and local keying

Media Encryption Key and the Access Key

This encryption method is as secure but much more flexible than encrypting directly with the access key. The access key can be changed without affecting the encrypted data, because the Media Encryption Key remains unchanged. If data were encrypted with the access key, as in the past, then changing the key would destroy data. Likewise, overwriting the Media Encryption Key does destroy data, resulting in an instantaneous cryptographic erasure of the entire drive.

If a SED is not configured with an access key, then data is readable as if the drive were not self-encrypting.

If a SED is configured with an access key, then the access key must be provided to unlock the drive, which remains unlocked only while powered. The drive locks itself upon losing power or shutting down, and the access key must be provided again.

This information also applies to partitions of a SED (called bands by the Trusted Computing Group [TCG]). Each partition has its own Media Encryption Key and optional access key (called a BandMaster by the TCG). AutoSED configures a small unsecured band for drive labels, followed by a single secured band spanning the rest of the drive. This access key is the key that is protected by AutoSED.

Threshold Secret Sharing and Local Keying

The AutoSED feature is a self-contained keying system, requiring no external Key Management Service (KMS). Exclusive to Dell, automatic local keying relies upon the concept of cryptographic secret sharing as discovered by Adi Shamir and specified in the Internet Draft Threshold Secret Sharing by David McGrew (draft-mcgrew-tss-03).

When a SED member is initially configured, AutoSED generates a new and unique access key. Every drive in the system is locked with this one key. Then, the Shamir algorithm is used to split the key into any number of pieces, called shares, which have the following properties:

1. For each set of shares, you can choose how many shares are needed to recover the key (for instance, 2-out-of-3 or 10-out-of-20). This number is the threshold.
2. Every time the key is split into a set of shares, the shares will be different even though the key stays the same. Shares can be combined only with shares from the same set; they are incompatible with shares from any other set.
3. Shares disclose no information about the key until the threshold is reached.

AutoSED always chooses to split the key such that one share is written to each active drive in the system (that is, non-spare, non-failed, non-foreign drives). The threshold is always half that number; more precisely, it is $(n+1)/2$. Therefore, the SEDset can automatically unlock itself whenever half of the drives are present. For the same reason, an adversary must possess half of the drives from the same SEDset to unlock it.

When AutoSED generates a backup set, this set consists of three shares with a threshold of two, which adds security and reliability to a sensitive process.

To destroy a set of shares, you could erase every share. However, if you erase only two shares from a backup set, the remaining share cannot recover the key and is useless.

Example: AutoSED Key Sharing

Consider an enclosure with 22 active drives and 2 spares:

1. When the array is first set up, AutoSED generates shares for all the active drives, which are in Set A. Because spares are excluded, Set A has 22 shares with a threshold of 11 (11-of-22). The array then generates and displays a backup, Set B, which is 2-of-3.
2. Remove a drive. Immediately, AutoSED destroys Set A, which means erasing the rest of Set A (the remaining 21 shares). AutoSED generates a new Set C containing the 21 active drives and having a threshold of 11-of-21. If another drive is removed, then AutoSED destroys Set C, and replaces it with Set D and a threshold of 10-of-20.
3. If both drives are removed by the same adversary, it now possesses one share from Set A and one share from Set C. However, each of those shares is useless by itself and also cannot be combined with each other. The adversary must simultaneously remove at least 11 drives to obtain enough shares from the same set.

Set B has remained usable. However, the same process applies:

1. Suppose one of the shares from Set B is compromised. Immediately, the administrator should destroy Set B, which means erasing the rest of Set B (the remaining 2 shares).
2. Then, the administrator should generate a new backup set, resulting in Set E. If a share from Set E is compromised, destroy Set E by erasing the remaining 2 shares, and generate Set F. The stolen share from Set B and the stolen share from Set E are useless individually and also cannot be combined. Also, they cannot be combined with the shares stolen from Set A and Set C.

Further reading:

Shamir, Adi, *How to Share a Secret*, Communications of the ACM, 22(11):612-613, 1979

McGrew, David, *Threshold Secret Sharing*, Network Working Group Internet Draft (draft-mcgrew-tss-03)

Trusted Computing Group, Enterprise SSC v1.0 Specification

20

About Monitoring

You can review comprehensive data about your array groups. Monitoring your PS Series array groups provides data that enables you to assess the health of your storage environment to quickly identify hardware and software problems.

Monitoring provides valuable data about the time and action of users on the system, protecting against security intrusions. Using Dell EqualLogic performance monitoring tools, you can gather information about group performance based on latency, IOPS, I/O rate, I/O size, and other data. With this data, you can quickly be informed of hardware, capacity, and performance-related problems, allowing you to make corrections and minimize downtime.

Monitored data is invaluable when accessing your customer support account and reporting problems directly to technical support.

- Events — Displays events to track operations and also to detect and solve problems before they affect performance or data availability. Audit messages are syslog events about administrator actions. They provide historical reference to actions such as logging in, logging out, creating a volume, setting up replication, and other events. In addition, you can use SNMP traps to track significant group events.
- Statistics — Displays information about the current and recent administrative sessions and the number of active iSCSI connections.
- Schedules — Displays the total number of snapshot and replication schedules.
- NAS Schedules — Displays NAS snapshot, replication, and data reduction schedules.
- Replication — Displays active replication activity (including inbound and outbound replication), and replication history.
- NAS Replication — Displays active NAS replication activity (including inbound and outbound replica containers), and NAS replication history.
- Sync Rep — Displays information about SyncRep volumes.
- Alarms and Operations — Displays a visual cue to alarm conditions in the group, as well as in-progress operations. Some operations might need administrator intervention. The cues are displayed for critical and warning alarms, actions, group operations, and failback operations.

Tools That Monitor and Manage Storage Performance

Dell EqualLogic provides the following tools for monitoring and managing the storage performance of your group:

- Group Manager GUI
- Group Manager Performance Monitor, available from the Group Manager GUI
- SAN Headquarters, which enables you to monitor multiple PS Series groups from a single GUI

About Using Group Manager to Monitor Performance Data

The Group Manager GUI provides dozens of views that provide comprehensive data about your array groups. From the Group Manager GUI, you can monitor the following data:

- Events
- Administrative sessions
- iSCSI connections
- Snapshot schedules

- Replication schedules, replication (including inbound and outbound replication), and replication partners
- Alarms and operations (including critical and warning alarms, actions, group operations, and failback operations), and storage pool free space
- Group members (including a specific member), the member health status, and member space
- Member enclosures, including power supplies and other hardware
- Control modules
- Disk drives
- Network hardware
- Volumes, collections, and snapshots, including current and requested status

About Displaying Performance Statistics with the Performance Monitor

The Group Manager Performance Monitor shows performance statistics for the drives or control modules in a member. The Performance Monitor GUI, accessible from the Group Manager GUI, shows statistical data at regular time intervals. You use icons to start or stop polling the data, and to navigate through the captured data. You can display up to four sets of statistics.

The Performance Monitor provides some flexibility, allowing you to change how data is displayed. You can view data as a chart, a histogram (bar chart), or a data table. For example, in chart view, you can click at any place along the graph and show details for a given time slice. You can also change the Performance Monitor display colors used in graphs, change the length of time between which data points are collected, and change the number of data points to save.

To more effectively monitor multiple PS Series groups from a single GUI, Dell recommends that you use the SAN Headquarters monitoring tool.

Start Performance Monitor from the Tools Menu

1. Open the Tools menu and click **Performance monitor**. The Group Manager Performance Monitor window opens.
2. Click **Add statistics** to open the Select Statistics dialog box.
3. Expand **Members**.
4. Expand a specific member.
5. Expand a component or statistics category.
6. Select a counter set to display.
7. Click **OK**.

In the Performance Monitor window, use the buttons and icons along the top of the window to control which data is displayed and how it is displayed.

Start Performance Monitor from the Navigation Pane Tree View

1. Click **Group**.
2. Expand **Members** and select a member from the list.
3. Either:
   - Click the **Disks** tab to monitor a drive.
   - Click the **Network** tab to monitor an Ethernet port.
4. Select the component that you want to monitor.
5. In the Activities panel, click the link for the type of statistics that you want to monitor.

In the Performance Monitor window, use the buttons and icons along the top of the window to control which data is displayed and how it is displayed.

NOTE: While it is possible to have several Performance Monitor windows open at the same time, Dell recommends running only one instance at a time to avoid possible unexpected results.

Performance Monitor Operation Icons

Table 54. Performance Monitor Operation Iconss shows the operation icons in the Performance Monitor window.

NOTE: The “Go to” icons work only when monitoring is paused.

Table 54. Performance Monitor Operation Iconss

Icon	Operation
▶	Start polling the data.
⏸	Stop polling the data.
⏮	Go to the start (first item).
❮	Go to the previous item.
❯	Go to the next item.
⏭	Go to the end (last item).

Add, Change, or Remove Statistics

You can display up to four sets of statistics in the Performance Monitor window.

To add more statistics:

1. Click **Add statistics** to open the Select Statistics dialog box.
2. Expand **Members**.
3. Expand a specific member.
4. Expand a component or statistics category.
5. Select a statistic to display.
6. Click **OK**.
7. To add more sets of statistics (up to four), repeat steps 1 through 6.

In the top-right corner of each statistics panel, use the two icons to either:

- Select a different set of statistics to display in that panel
- Close (remove) that panel

  When a panel is removed, all other panels remain open in the Performance Monitor window and are resized to fill the window. You can open a new panel in the window (up to four total).

Using Performance Monitor Counter Sets

Counter sets in the Performance Monitor display detailed usage statistics about SAN components. You can create custom counter sets or use preconfigured sets that the system provides. You can also modify and delete custom counter sets.

NOTE: You cannot modify or delete preconfigured counter sets.

To create (add) or modify counter sets:

1. In the Performance Monitor window, click **Add statistics**.
2. In the Select Statistics dialog box, click **Counter sets**.
3. In the Counter Set Management dialog box, select the type of statistics that you want to include in the counter set.
4. Click the appropriate link to open the Counter Set Modification dialog box.
5. Type a name and description, and then select the statistics that you want to include in the counter set.
6. Click **OK** in the Counter Set Modification dialog box.
7. Click **Close** in the Counter Set Management dialog box.

To delete counter sets:

1. In the Performance Monitor window, click **Add statistics**.
2. In the Select Statistics dialog box, click **Counter sets**.
3. In the Counter Set Management dialog box, select the counter set that you want to delete.
4. Click the **Delete** link.
5. In the Delete Counter Set Confirmation dialog box, click **Yes**.

Change How Data Is Displayed

Table 55. Changing How Data Is Displayed shows the icons that you use to change the data display. You can view the data as a chart, a histogram (bar graph), or a data table.

Table 55. Changing How Data Is Displayed

Icon	Data Format	Options
📈	Chart (line graph)	▦ Displays data on a linear scale ▦ Displays data on a logarithmic scale ↕ Resizes the display (scales in or out to fit)
📊	Histogram (bar graph)	
▦	Data table	None

NOTE: The display mode that you select applies to all panels currently open in the Performance Monitor window.

Display Data for a Specific Point in Time

In the chart view (for both linear and logarithmic scales), when you click inside the data window, the cursor changes to a crosshair. You can click at any place along the graph to display the details for that time slice.

For example, you can click the crosshairs on the peak of a graph line to display the data for that moment in time in the panel on the left (see Figure 35. Performance Monitor – Select Data Point). In this example, at the peak of the graph, the group is processing 303 output requests.

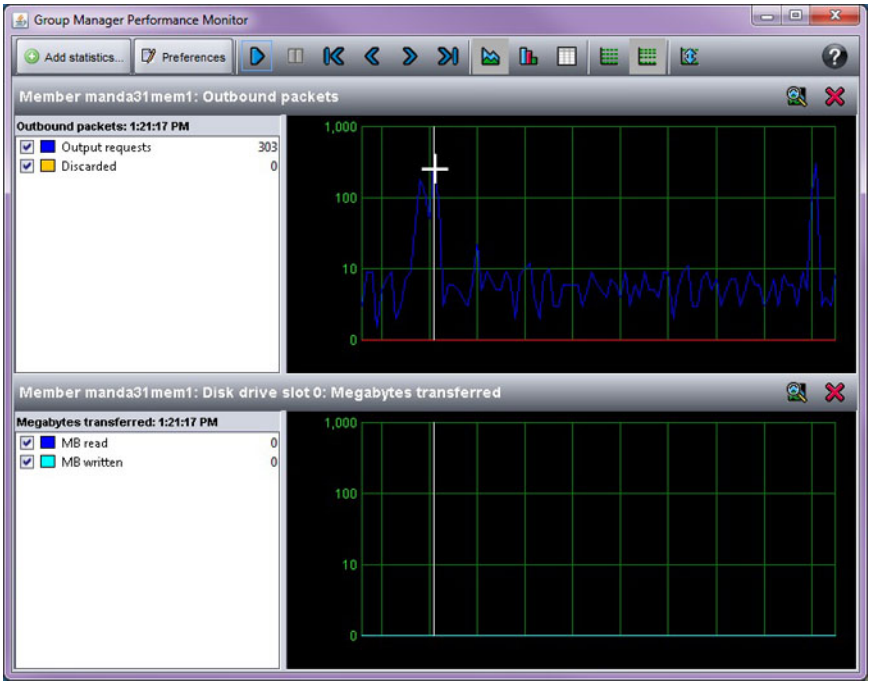
Group Manager Performance Monitor
Add statistics...
Preferences
Member manda31mem1: Outbound packets
1,000
100
10
0
Member manda31mem1: Disk drive slot 0: Megabytes transferred
1,000
100
10
0

Outbound packets: 1:21:17 PM

☑ 🟦 Output requests 303
☑ 🟨 Discarded 0

Megabytes transferred: 1:21:17 PM

☑ 🟦 MB read 0
☑ 🟦 MB written 0

Figure 35. Performance Monitor – Select Data Point

Customizing the Performance Monitor

Within the Performance Monitor window, you can change the following items:

- Colors used in graphs

  See Change the Display Colors.
- Values for data points:
  - Length of time between which data points are collected
  - Number of data points to save

  See Change the Data Collection Values.

Change the Display Colors

On the left side of the window within each statistics panel, you can change the colors used in the Performance Monitor display.

1. Click the colored box (for example, 🟩) to open the **Select a color** dialog box (see Figure 36. Performance Monitor – Select a Color Dialog Box). You can either:
   - Choose from a predefined swatch panel
   - Use one of the other tabs to specify a custom color value
2. Click **OK** when you have finished.

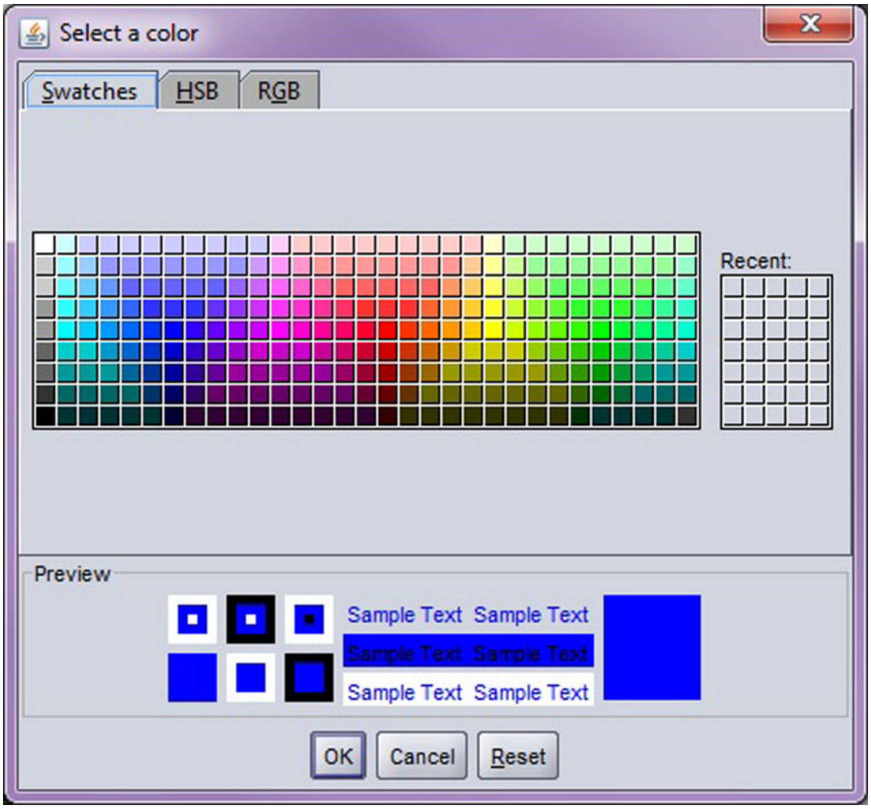

Figure 36. Performance Monitor – Select a Color Dialog Box

Change the Data Collection Parameter Values

You can change the following parameter values for data collection as needed:

- Time between data points
- Number of data points to save

Table 56. Data Collection Values shows the parameters and their default and maximum values.

Table 56. Data Collection Values

Parameter	Default	Maximum	Option
Time between data points	1 second	60 seconds	Predefined intervals (1, 5, 10, 30, and 60) or any integer value between 1 and 60
Number of data points to save	100 points	1000 points	Predefined values (100, 250, 500, or 1000) or enter any integer value between 100 and 1000

To change the values:

1. In the Performance Monitor window, click **Preferences**.
2. In the dialog box, either select one of the predefined values from the list or specify another value.
3. Click **OK**.

As you choose different values for the parameters, the dialog box shows you the total amount of time that you can save with that combination of values. For example, using an interval of 60 seconds and saving 300 data points, you can store up to 5 hours of statistics.

The changes take effect immediately. You do not have to stop and restart polling.

About SAN Headquarters

SAN Headquarters (SAN HQ) is a performance-monitoring tool that enables you to monitor multiple PS Series groups from a single graphical user interface (GUI).

SAN HQ gathers and formats configuration and performance data into easy-to-view charts and graphs. Analyzing this data can help you improve performance and more effectively allocate group resources.

SAN HQ allows to you deploy Dell SupportAssist for diagnostic data collection on a weekly basis, on-demand as needed, or when critical events occur. Diagnostic data collections are automatically uploaded to Dell Technical Support for analysis, unless you disable automatic uploads. In that case, the data is encrypted and stored locally.

By default, Group Manager displays an Action alarm in the Alarms and Operations panel in the lower part of the window whenever SAN HQ is not monitoring the group.

Using SAN HQ, you can:

- Quickly be informed of problems related to hardware, capacity, and performance
- Improve performance by identifying performance bottlenecks
- Obtain comprehensive information about group performance based on latency, IOPS, I/O rate, I/O size, and other data. Using 95th percentile reporting, remove the top 5 percent of spikes in data for a more accurate picture of your storage performance.
- View real-time data for group member or volume I/O and save the results for future analysis
- From a single SAN HQ client, monitor group performance data from multiple servers
- Determine how the group is performing, relative to a typical I/O workload of small, random I/O operations. This information can help you determine if a group has reached its full capabilities, or whether you can increase the group workload with no impact on performance.
- Display performance-based load-balancing data for the group (the default), for all pools on the group, or for all members in the group.
- For a selected group, pool, or member, apply different RAID policies and analyze the performance benefits
- View events, audits, and group alerts
- More effectively allocate group resources by identifying underutilized resources
- Enable multiple individuals to access and monitor the same performance data
- Preserve group performance data for later analysis by creating archives
- Create customized reports of group performance data
- Export group performance data to a spreadsheet
- Specify favorite views
- Visualize synchronous replication volumes and NAS containers
- Display space consumed by recoverable volumes in the recovery bin
- Display space available for snapshot borrowing
- See the number of Ethernet ports with active and inactive data center bridging (DCB), and the number of ports incompatible with DCB

For complete information about installing and using SAN Headquarters, see the *Dell EqualLogic SAN Headquarters Installation and User's Guide*.

Manage SAN Headquarters Action Alarm

In Group Manager, the SAN Headquarters (SAN HQ) section of the Group Configuration General panel lists servers from which SAN HQ is currently monitoring or had previously monitored the group.

By default, if no SAN HQ server is currently monitoring the group, the Alarms and Operations panel at the bottom of the Group Manager GUI displays an Action alarm.

To hide the alarm and configure Group Manager to never display a SAN HQ monitoring alarm in the future:

- Click the **Hide SAN HQ reminder** option from within the alarm.

To reverse this configuration and reenable the display of SAN HQ monitoring alarms:

1. Click **Group** → **Group Configuration** → **General**.
2. In the SAN Headquarters section of the panel, select **Enable reminder in Alarms panel**.

Monitor Group Members

Member hardware problems typically cause event messages and alarms. Monitor the member hardware and replace any failed components immediately.

To display member information:

- Click **Group** → **Group** ***group_name***

The **Group Information** panel shows the general settings of the group, and the volumes, snapshots, and collections that are associated with that group.

The **Group Disk Space** panel shows the total amount of free space in the group and in each pool (if applicable).

The **Storage Pools and Group Members** panel lists all the members, the pool each member belongs to, their capacity and amount of free space, RAID policy, number of disks, member status, health status, and whether disk encryption is being used.

Check the following items:

- Member status – If a member is offline, investigate the cause. Volumes with data on an offline member are also offline. If a member has a problem, double-click the member to display additional information.
- Low free space – Low free space in a member might indicate that overall group space is low. You can free space in a member by adding more members to the same pool (the group distributes volume data across the pool members).

About Monitoring Enclosure Elements

Enclosure elements include power supplies, cooling modules and fans, EIP and OPS cards, and channel cards.

Power Supplies

- A member can survive one power supply failure. Replace failed power supplies as soon as possible.
- For proper cooling, do not remove a power supply until you have a replacement.
- For information about replacing a power supply, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider.
- The Power Supplies panel shows the status of the power supplies. The number and type of hardware components shown depends on your array model.

Cooling Modules and Fans

- A member has two or three cooling modules and multiple fans. Most PS Series arrays use power supplies that have integrated cooling modules.
- Periodically, check the room temperature where the hardware is located and make sure that the room is sufficiently cool and ventilated. Also, make sure the fan trays and cooling modules have no red LEDs, and monitor the member temperature.
- A member can survive one cooling module failure. Replace failed cooling modules as soon as possible.
- Multiple fan failures increase the array temperature. A high temperature results in event messages. The array might shut down before damage occurs.
- Some PS Series arrays also show the ambient temperature, which is calculated in Celsius from the two sensor temperatures with the highest temperatures, using the following formula:

  ```
  ((Backplane Sensor 0 + Backplane Sensor 1) / 2) – 7
  ```

EIP or OPS Card

- Some array models include an Enclosure Interface Processor (EIP) card, and others contain an OPS (operations) card. An array continues to operate if the EIP or OPS card fails. You can replace the failed EIP or OPS card with no impact on group operation.
- In the Member Enclosure window, the EIP card panel shows the EIP card status. The OPS card panel shows the OPS card status.

Channel Cards

- Some array models include redundant channel cards. An array continues to operate if a channel card fails. You can replace the failed channel card with no impact on group operation.
- For information about replacing channel cards, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider.

Monitor a Specific Member

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Status** tab.

The General Member Information, Member Health Status, and Member Space panels display information about the selected member.

Use LEDs to Identify a Member

If a hardware failure occurs in a member, LEDs on the group member containing the failed component will illuminate.

To help you identify a member, you can also illuminate (flash) these LEDs:

- Fan tray LED
- Control module ERR LED on the member chassis

To illuminate a member’s LEDs:

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Maintenance** tab.
4. Click **Start LEDs flashing** in the Identification panel.

To return the member’s LEDs to normal operation:

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Maintenance** tab.
4. Click **Stop LEDs flashing**.

CAUTION: You must shut down the member before you turn off power.

Monitor the Member Enclosure

The member enclosure information includes the power supplies, cooling fans (usually integrated into the power supplies), and, on some array models, channel cards and an EIP or OPS card, depending on the array model.

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Enclosure** tab.

The various panels display information about the selected member.

Monitor Control Modules

Each group member has one or two control modules installed. One control module is designated as active (responsible for serving I/O to the member). On the active control module the LED labeled ACT is lit.

In a dual control module array, the other control module is secondary (mirrors cache data from the active control module). On startup, either control module can be designated active or secondary, regardless of its previous status.

Under normal operation, the status of a control module (active or secondary) does not change unless you restart the member.

In a single control module array, if the control module fails, the member is offline.

In a dual control module array, if the active control module fails, the secondary control module becomes active and begins serving I/O (called control module failover). I/O should continue if you connect cables to the newly active control module.

For information about replacing control modules, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider.

To display control module information:

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Controllers** tab.

Each Control Module Slot panel shows the following information.

NOTE: In some cases, after a control module has been upgraded or replaced, the Group Manager GUI might display a failed status for the control module's power supply or battery. The GUI will update automatically after a few minutes to reflect the correct status.

- Status
- Boot time
- Cache battery status or cache-to-flash module and NVRAM battery status, depending on the array model
- Model number and type
- Boot ROM version
- PS Series firmware version

An empty slot means that a control module is not installed or has failed. The system might temporarily indicate that a slot is empty while a control module is rebooting.

For information about replacing a control module, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider. Do not remove a failed control module until you have a replacement.

The Memory Cache panel displays the cache mode. Control module and battery status affect the cache mode. Write-through mode might impair performance. Identify why the cache is in write-through mode and correct the problem, if necessary.

Control Module Status

Table 57. Control Module Status describes status values for control modules.

Table 57. Control Module Status

Status	Description	Solution
`active`	Serving I/O to the member	None needed; informational
`secondary`	Mirroring cache data from the active control module	None needed; informational

Cache Battery Status

Table 58. Cache Battery or Cache-to-Flash Module Status describes status values for control module cache batteries or cache-to-flash modules, depending on the array model, and provides solutions for any problems.

Table 58. Cache Battery or Cache-to-Flash Module Status

Status	Description	Solution
`ok good`	Battery or cache-to-flash module is fully charged.	None needed; informational.
`failed`	Battery or cache-to-flash module has failed.	Contact your service provider for more information.
`missing battery`	Battery is missing.	Contact your service provider for information. Does not apply to the cache-to-flash module.
`low voltage`	Battery is below the limit for normal operation.	If the battery status is `low voltage` for an extended period of time, contact your PS Series service provider for more information. Does not apply to the cache-to-flash module.
`low voltage, is charging`	Battery is charging but is still below the limit for normal operation.	If the battery status is `low voltage, is charging` for an extended period of time, contact your PS Series service provider for more information. Does not apply to the cache-to-flash module.
`good battery, is charging`	Battery is charging but has enough charge for normal operation.	None needed; informational. Does not apply to the cache-to-flash module.

Channel Card Status

Table 59. Channel Card Status describes status values for channel cards and provides solutions for any problems.

Table 59. Channel Card Status

Status	Description	Solution
`good`	Channel card is functioning normally	None needed; informational.
`failed`	Channel card failure	Contact your PS Series support provider for information about replacing a channel card.
`not-present`	Channel card is missing or status is unavailable	Contact your PS Series support provider for information about installing or replacing a channel card.
`unknown`	System cannot determine the channel card status	Contact your PS Series support provider.

For information about replacing channel cards, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider.

NVRAM Battery Status

Table 60. NVRAM Battery Status describes status values for control module NVRAM coin cell batteries and provides solutions for any problems.

NOTE: Some arrays do not have an NVRAM battery.

Table 60. NVRAM Battery Status

Status	Description	Solution
`good`	Battery installed and fully charged	None needed; informational.
`bad`	Battery failure	Contact your PS Series support provider.
`not-present`	Battery is not installed	
`unknown`	Battery status is not known	

Monitor Disk Drives

Make sure you detect and replace failed drives as soon as possible. Although spare disks and RAID protect data against drive failures, multiple failures might put data in jeopardy.

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Disks** tab.

The Disk Array Summary panel shows the drives in the member. The number and type of drives shown depends on your array model.

The Installed Disks panel shows more information about each drive, including the slot, type, model and revision, size, status, and errors. Closely monitor drives with errors.

CAUTION: Never pull out and then immediately reinsert a drive.

Disk Drive Status

Table 61. Disk Drive Status describes status values for drives and provides solutions for any problems.

NOTE: When you remove a drive from an array, Dell recommends that you not reinsert it immediately. Instead, wait for the system to recognize that the drive is missing before reinserting it.

Table 61. Disk Drive Status

Status	Description	Solution
`small`	Drive is smaller than other drives in the member. The drive cannot be used in the member.	Replace the drive with one that is the same capacity or greater than the installed drives.
`failed`	Drive failure	Contact your PS Series support provider for information about replacing failed drives.
`foreign label`	Drive has a foreign label. The drive was probably removed from a different array and then installed in this array.	To use the drive, click **foreign disk** and clear the label.
`history of failures`	Previously failed drive	Contact your PS Series support provider. To use the drive, click **history of failure** and agree to use the drive.
`online`	Drive is functioning.	None needed; informational

Status	Description	Solution
`copying to spare`	Data is being written to a spare drive.	None needed; informational
`unsupported version`	Drive is running an unsupported firmware version.	Contact your PS Series support provider.

When a drive in a RAID set fails, a member behaves as follows:

- If a spare drive is available — Data from the failed drive is reconstructed on the spare. During the reconstruction, the RAID set that contains the failed drive is temporarily degraded.
- If a spare drive is not available, and the RAID set has not reached the maximum number of drive failures — The RAID set that contains the failed drive is degraded. For RAID 5, RAID 50, or RAID 6, performance might decrease.

  ⚠ **CAUTION: A drive failure in a RAID 5 or RAID 10 set that is degraded might result in data loss.**
- If a spare drive is not available, and the RAID set has reached the maximum number of drive failures — The member is set offline, and any volumes and snapshots that have data stored on the member are set offline. Data might be lost and must be recovered from a backup or replica.

When you replace a failed drive, a member behaves as follows:

- If a spare drive was used — The new drive automatically becomes a spare, with no effect on performance.
- If a RAID set was degraded — Data is automatically reconstructed on the new drive and performance goes back to normal after reconstruction.
- If a member was offline because of multiple RAID set drive failures — Any volume snapshots with data on the member are set offline and data might be lost.

In some cases, a member might detect a problem with a drive. The member automatically copies the data on the failing drive to a spare drive, with no impact on availability and little impact on performance. The group generates event messages informing you of the progress of the copy-to-spare operation. I/O is written to both drives until the copy-to-spare operation completes. If the drive completely fails during the operation, data is reconstructed on the spare using parity data as usual.

Replace any failed drives immediately. For information about replacing drives, see the *Hardware Owner's Manual* for your array model or contact your PS Series support provider.

Monitor Network Hardware

A member must have at least one functioning network interface connected to a network and configured with an IP address. Each control module has multiple Ethernet ports.

If you experience network problems, group members might lose the ability to communicate with each other over the network. In such a group, some management operations are not allowed. For example, you cannot change the IP addresses of an isolated member.

If the members of a group cannot communicate, identify and correct the network problems. Correcting these problems restores the group to normal full operation, including network communication.

To display the network information:

1. Click **Group**.
2. Expand **Members** and select the member name.
3. Click the **Network** tab to display the IP Configuration panel.

The IP Configuration panel provides the following information:

- Current status (under **Status**) — Current status of the network interface:
  - up — Operational, connected to a functioning network, configured with an IP address and subnet mask, and enabled

  - down — Not operational, not connected to a functioning network, not configured with an IP address or subnet mask, or disabled
- Port failover status (under **Controller**) — Current status of the controller:
  - Primary — no vertical port failover
  - Secondary — vertical port failover has occurred
  - Unknown — value cannot be determined
- Requested status (under **Network interface**) — Status set by administrative action:
  - enabled — Configured and serving I/O
  - disabled — Not serving I/O, but might be configured

If the current status is `down` but the requested status is `enabled`, identify and correct the error. For example, under **Network interface**, check the following settings:

- Speed — Make sure that the interface speed is adequate.
- Packet errors — A few packet errors are not usually a problem. If a large number of packet errors occur, a network problem or a network interface or port failure might exist. Identify and correct the problem.

  To protect against network interface or port failure, connect multiple network interfaces on both control modules to the network.

Monitor iSCSI Connections to a Member

To display all connections to a member:

1. Click **Group**.
2. Expand **Members** and then select the member name.
3. Click the **Connections** tab.

The iSCSI Connections panel shows information about the initiator address, which volume or snapshot it is connected to (**Target** column), how long the connection has been active, and which Ethernet port the initiator is using.

Check for multiple initiators writing to the same iSCSI target. This situation can cause target corruption if not handled correctly by the servers.

Monitor iSCSI Connections

Check for multiple initiators writing to the same target. This situation can cause volume corruption if not handled correctly by the servers.

To monitor iSCSI connection statistics for all the targets (volumes and snapshots) in the group:

- Click **Monitoring → iSCSI Connections**.

About Storage Performance

To optimize your SAN performance, be sure to regularly analyze your environment. If you build these actions into your schedule, you will prevent performance issues:

- Check for damaged hardware

  Eliminate bad hardware as the initial cause of performance problems. The issue might be as simple as an unplugged power cable.
- Check the volume I/O latencies

  One of the leading indicators of the health of your SAN is latency, the time from the receipt of the I/O request to the time the I/O is returned to the server. Many applications exhibit significant performance degradation when latencies are consistently above 50 ms. To know whether your storage environment is performing optimally you must understand how your applications function with the PS Series SAN.

- Check your PS Series group capacity

  A key component of the health of your PS Series group is capacity. To fully understand the capacity for new applications or support the growth of existing servers, you must examine the overall group and pool capacity, storage utilization statistics, thin-provisioned space, and space used for replication. To ensure a healthy SAN, it is important to detect any sudden or unexpected changes in capacity utilization.
- Determine overloaded SAN resources

  When high latency cannot be attributed to incorrect configuration of the storage environment, the SAN resources might be overloaded. Use SAN Headquarters to help make this determination. SAN Headquarters can identify random IOPS, network performance, storage pool capacity, iSCSI connections, MPIO connections, and other indicators.

Figure 37. Process for Analyzing Data shows the process for analyzing data.

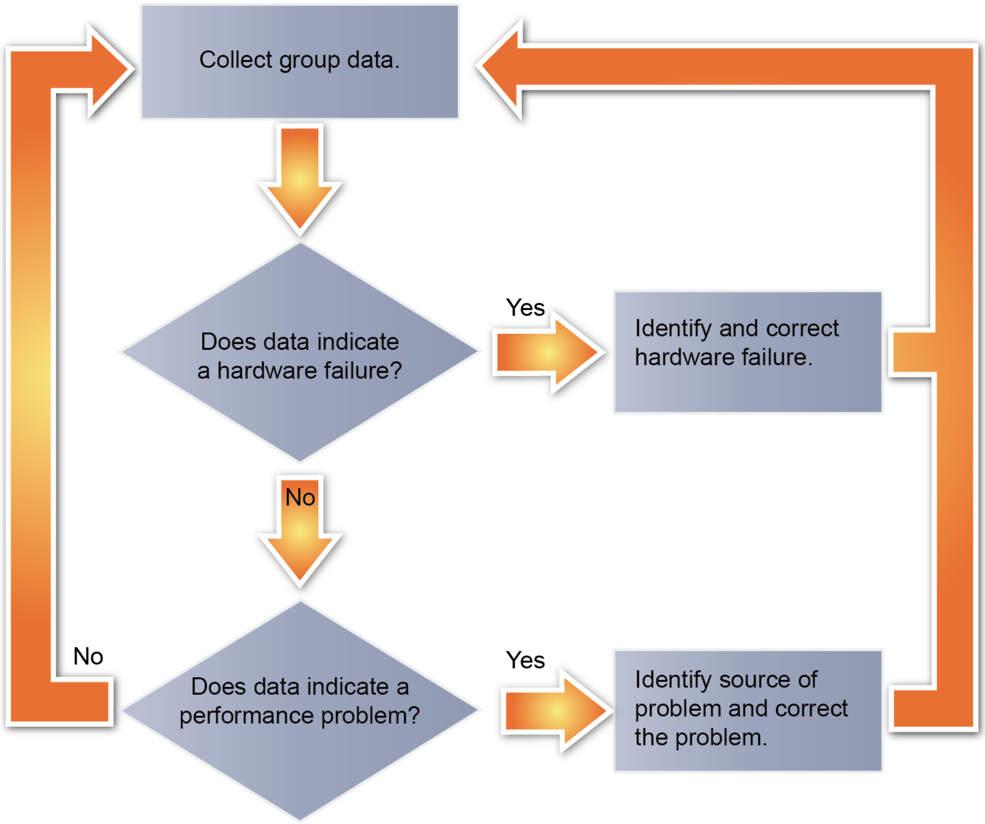
Collect group data.
Does data indicate a hardware failure?
Yes
Identify and correct hardware failure.
No
Does data indicate a performance problem?
Yes
Identify source of problem and correct the problem.
No

Figure 37. Process for Analyzing Data

To analyze your data, you should also understand the following concepts:

- How group data is compressed and therefore how it might be less precise as it ages. For information, see the *Dell EqualLogic SAN Headquarters Installation and User's Guide*.
- Basic performance terminology and the type of data collected.
- The areas in your environment that can be sources of performance problems and which areas the monitoring tools look at.
- How your applications use group storage resources.

Best Practices for Managing Storage Performance

As a general best practice, focus on issues elevated from your user environment. Users might have the following complaints:

- Response time is too slow.
- An action takes too long to complete.
- An action did not finish on time.

Follow these guidelines to analyze the performance problem:

- Are users getting the response time they expect? If not, identify which area might be causing the problem:
  - Operating system problem, as it interacts with storage
  - Network problem
  - Application being run or accessed
  - Storage environment
- Use the 80/20 rule. By focusing on 20 percent of the most likely causes of a performance issue, you will solve 80 percent of the problems.
- Keep the host perspective in mind when managing the arrays. In particular:
  - Make sure time is synchronized on all monitoring areas (host, array).
  - Never assume the cause of a problem. Other issues might be causing skewed data.

General best practices for solving performance problems include:

- Fix hardware problems immediately, even if the array that is down is the redundant array. Replace failed disks or failed control modules.
- Always validate the data you are analyzing to ensure it is accurate. For example, you might need to inspect the physical array to see if it is powered on or check if your switches are connected properly to handle MPIO.
- Consider the size of your installation when analyzing your storage data. For example, an enterprise-level installation with a large volume of users and data will experience a greater impact from a small degradation of IOPS than a small company with few users.
- When setting up email notification, determine what types of information are most useful. Streamline the information you receive as much as possible. If email notifications are set too broadly, an actual problem might be obscured by too much information.
- When using SAN Headquarters to monitor your groups, remember that the historical data collected degrades over time. Use more current data for your analysis.
- Look at event and audit logs for other issues.

Common Hardware Issues

Identifying hardware performance issues can eliminate additional effort elsewhere. Hardware failures can also be a source of performance problems. In addition, the combination of hardware and firmware can affect performance, as can various disk types with different performance characteristics.

The basic steps for solving any IT problem also apply to the SAN. Table 62. Hardware Issues Affecting SAN Performance lists some common problems that you should watch for and correct immediately.

Table 62. Hardware Issues Affecting SAN Performance

Damaged Hardware	Typical Symptom	Detected By	Possible Corrective Actions
Server NIC	Malformed packets	Monitor errors at switch	Update NIC drivers Replace NIC
Bad cable Wrong class of cables	Visible damage Malformed packets	Visual inspection Monitor errors at switch	Replace cable
Defective switch	Spontaneous restarts Random lockup	Monitor switch with appropriate network	Update switch firmware Replace switch
Defective array hardware	Alerts	Monitor PS Series group Monitor SAN Headquarters	Contact Dell customer support to replace malfunctioning component

Damaged Hardware	Typical Symptom	Detected By	Possible Corrective Actions
		Set up email alerts on group and SAN Headquarters	

As best practice, use the SAN Headquarters GUI to help identify hardware-related issues. SAN Headquarters easily tracks the array model, service tag, and serial number, plus RAID status and policy, and firmware version. In particular, SAN Headquarters provides information about:

- Hardware alerts

  The SAN Headquarters Alerts panel shows hardware problems that might affect performance, such as a failed disk or a network connection that is not Gigabit Ethernet.
- Network retransmissions

  A sustained high TCP retransmit rate (greater than 1 percent) might indicate a network hardware failure, insufficient server resources, or insufficient network bandwidth.
- RAID status

  A degraded, reconstructing, or verifying RAID set might adversely affect performance. In some cases, performance might return to normal when an operation completes.
- Low pool capacity

  Make sure free space in each pool does not fall below the following level (whichever is smaller):

  – 5 percent of pool capacity
  – 100GB times the number of pool members

  Otherwise, load-balancing, member-removal, and replication operations do not perform optimally. Low free space also negatively affects the performance of thin-provisioned volumes.

About Analyzing SAN

If you are sure that no hardware problems exist, it is best practice to use SAN Headquarters to review performance statistics to identify other potential problems. These statistics provide a good indication of overall group performance and might help you identify areas where performance can be optimized.

The following statistics provide common indicators of performance problems: I/O latency, I/O load, IOPS, I/O size, network load, network rate, and queue depth.

Average I/O Latency

One of the leading indicators of a healthy SAN is latency. Latency is the time from the receipt of the I/O request to the time that the I/O is returned to the server.

Latency must be considered along with the average I/O size, because large I/O operations take longer to process than small I/O operations.

The following guidelines apply to I/O operations with an average size of 16KB or less:

- Less than 20 ms — In general, average latencies of less than 20 ms are acceptable.
- 20 ms to 50 ms — Sustained average latencies between 20 ms and 50 ms should be monitored closely. You might want to reduce the workload or add additional storage resources to handle the load.
- 51 ms to 80 ms — Sustained average latencies between 51 ms and 80 ms should be monitored closely. Applications might experience problems and noticeable delays. You might want to reduce the workload or add additional storage resources to handle the load.
- Greater than 80 ms — An average latency of more than 80 ms indicates a problem, especially if this value is sustained over time. Most enterprise applications will experience problems if latencies exceed 100 ms. You should reduce the workload or add additional storage resources to handle the load.

If the average I/O operation size is greater than 16KB, these latency guidelines might not apply. If latency statistics indicate a performance problem, examine the total IOPS in the pools. The storage array configuration (disk drives and RAID level) determines the maximum number of random IOPS that can be sustained. EqualLogic customer support or your channel partner can help size storage configurations for specific workloads.

Also, review the latency on your servers. If the storage does not show a high latency but the server does, the source of the problem might be the server or network infrastructure. Consult your operating system, server, or switch vendor for appropriate actions to take.

Estimated I/O Load

Based on latencies, IOPS, hardware, and the RAID configuration, the estimated I/O load is relative to the theoretical maximum capability of the group, pool, or member. Because the load value is an estimate, use it only as a general indicator. The I/O load can be:

- Low — Minimal load. Latencies are low.
- Medium — Typical load. Usually, either member throughput is greater than 1MB/second or IOPS are greater than 50, and average latencies are above 20 ms. If this load is sustained, you should monitor the group carefully.
- High — Load that is approaching the theoretical maximum capability. Usually, member throughput will be greater than 1MB/second or IOPS will be greater than 50, and average latencies will be above 50 ms. If this load is sustained, you should investigate further.
- Unknown — Indicates that SAN Headquarters cannot currently calculate the I/O load because of invalid SNMP counters or because of the group workload.

Average IOPS

IOPS are a good way to measure overall I/O activity and how much work is being done, especially if you also consider the average I/O size. However, IOPS do not indicate whether the storage system is overloaded or operating within its limits.

Average I/O Size

In general, the average size of an I/O operation is 16KB or less. The larger the I/O operation size, the longer it takes to process, which might affect latencies. Also, large I/O operations usually reduce the total number of IOPS.

I/O size can be useful in understanding workload characteristics, especially when measured at the volume level.

Network Load

The network load is the percentage of the theoretical maximum network bandwidth that is being used for sending I/O or receiving I/O, whichever has the highest value. The theoretical maximum bandwidth is based on the negotiated link speed for all active network interfaces. The network load percentage provides a quick measure of network activity. Network ports are typically full-duplex.

The network is rarely a bottleneck in a SAN. Usually, network bandwidth is underutilized, especially with random I/O workloads. However, bandwidth might become fully utilized during highly sequential operations, such as a backup; in most cases, fully utilized bandwith does not indicate a problem with the system.

Using all the available bandwidth on one or more member Ethernet interfaces generates an alert.

Network Rate

In general, the network rate should be 100 percent to 200 percent of the I/O (iSCSI) traffic. A network rate that is significantly more than 200 percent might indicate a problem.

Queue Depth

SAN Headquarters displays the queue depth (average number of outstanding I/O operations at the start of each incoming I/O operation) for each disk drive (raw I/O), volumes (only iSCSI traffic), groups, and pools. A queue depth of zero indicates no I/O activity. High or sustained queue depths might indicate that the group is under a high load.

NOTE: A group must be running PS Series firmware version 5.0 or a later version to display iSCSI queue depth for a volume.

Monitor Administrative Sessions

To monitor administrative statistics:

1. Click **Monitoring**.
2. Select **Administrative Sessions**.

The **Active Sessions** and **Most Recent Login by Account** panels display information about the different sessions.

Monitor Snapshot Schedules

To monitor snapshot schedules:

1. Click **Monitoring**.
2. Below Schedules, select **Snapshot Schedules**.

To see more detail about a schedule, move the pointer over a schedule entry in the panel. A pop-up window opens, and shows this additional information:

- Schedule type (once, daily, hourly)
- Next time the schedule runs
- Number of snapshots to keep

You can take the following actions on a schedule:

- To modify a schedule, either double-click the schedule, or select it and click **Modify**. The **Modify schedule** dialog box opens. Make the changes, then click **OK**. You can change the following items:
  - Schedule name
  - Run date and time
  - Number of snapshots to keep
- To disable or enable a schedule, select it, right-click in the highlighted area, and then select the appropriate option from the menu. If you are disabling a schedule, click **Yes** in the confirmation dialog box.
- To delete a schedule, select it and click the **Delete**. In the confirmation dialog box, click **Yes** to delete the schedule.

Monitor Volumes and Snapshots

Use the following procedures to monitor volume and snapshot data at the group, volume, or snapshot level, respectively.
To display information for all volumes in a group:

1. Click **Volumes** → **Volumes**.
2. In the **Volumes** panel, check the status and attributes of every volume in the group.

To display detailed information about a specific volume:

1. Click **Volumes**.
2. Expand **Volumes**.
3. Select the volume name that you want to monitor.
4. Click the tabs to display specific information about the volume.

To display snapshot information:

1. Click **Volumes**.
2. Expand **Volumes**.

3. Expand a volume name.
4. Select a snapshot timestamp.
5. Click the tabs to display specific information about the snapshot.

Information That You Should Monitor

- Check the status of all volumes and snapshots. Each volume and snapshot has the following status values:
  - Current status — Actual status of the volume or snapshot, regardless of the requested status.
  - Requested status — Administrator-applied setting for the volume or snapshot. For example, an administrator can set a volume online or offline.

  For volumes, the current status is displayed in the Status column of the Volumes panel (**Volumes** → **Volumes**). To see the requested status, move the pointer over the volume information to display details in a pop-up window.

  For snapshots, the current status is displayed in the Status column of the Snapshots panel (**Volumes** → **Volumes** → **selected volume** → **Snapshots tab**). To see the requested status, move the pointer over the snapshot information to display details in a pop-up window.

  Under normal conditions, the requested status and current status are the same. However, an event in the group can result in a current status that differs from the requested status. For example, if an administrator sets a volume online, but a member containing volume data is shut down, the requested status is `online`, but the current status is `offline (member down)`. Always investigate when the current status is different from the requested status.

  The group sets a volume or snapshot offline if network problems occur or if a member that contains volume data is not available. If the group sets a volume offline because of a problem, the group also sets all its snapshots offline.

  The group also sets a thin-provisioned volume offline if the volume's maximum in-use space setting is less than 100 percent and a write exceeds this value.
- Check the snapshot reserve free space and the space recovery setting. If free snapshot reserve is exceeded, the space recovery setting controls whether the oldest snapshots are deleted to create free space for new snapshots, or the volume and its snapshots are set offline.

  You can also increase the snapshot reserve or change the space recovery setting. If data reduction has been enabled on the volume, snapshot reserve is permanently disabled.
- For volumes that are not thin-provisioned, check for in-use space that is near a volume's reported size. When free volume space is exhausted, writes to the volume fail, potentially disrupting applications, but the volume remains online.

  To increase free volume space, increase the reported volume size.
- For thin-provisioned volumes, check for in-use space that is near the maximum in-use space setting.

  Thin-provisioned volumes are set offline if their maximum in-use space is set to less than 100 percent and a write exceeds this value. If the maximum in-use space is set to 100 percent, and a write exceeds this value, the write fails, but the volume remains online.

  To increase free volume space, increase the reported volume size. You can also increase free volume space by increasing the value of the maximum in-use space value up to 100 percent.

About Monitoring Replication

If you are replicating volume data, you should monitor replication operations to ensure that each operation completes.

For example, you should monitor information about:

- Outbound replication (all volumes on the group configured for replication)
- Inbound replication (all replica sets stored in the group from all partners replicating to this group)
- Replication history (history of all outbound replications)

NOTE: Replication history displays the last 10 replicas only.

In addition, you should monitor the usage of delegated space. If free delegated space is not available, replica reserve cannot increase automatically. You can also monitor replica reserve for a volume. Insufficient replica reserve limits the number of replicas.

Table 63. Best Practices for Monitoring Replication suggests areas that you should monitor when performing replication between groups, and describes best practices to avoid replication issues associated with these areas.

Table 63. Best Practices for Monitoring Replication

Issue	Best practice to avoid issue
Incomplete replication operations	If a replication operation fails to complete, you might need to increase replication space.
Incomplete manual transfer operations or failback operations	Some operations require multiple tasks that administrators must complete. Make sure you complete all multitask operations.
Low free delegated space	If delegated space is low, replica reserve space might not be able to increase automatically to store new replicas.
Number of replicas	If too many replicas exist, consider decreasing the replica reserve percentage. If too few replicas exist, consider increasing the replica reserve percentage. A low free replica reserve can indicate optimal use of replica reserve space, if the desired number of replicas exist.

Monitoring Replication History

You should periodically monitor information about past replication operations, including volume and partner information, start time, end time, and duration for the replication operation, and the amount and speed of data transferred.

Also, you should examine the replication-duration information. If you see long replication times, make sure the network connection between the partners is sufficient. A slow network link between the partners can cause long replication times. If a replication operation makes no progress, the group generates a warning event. Make sure you have adequate network bandwidth between the groups, in addition to full IP routing. If necessary, increase the network bandwidth.

Additionally, check how much data you are replicating. You might want to use manual transfer replication if you are transferring a large amount of data.

Monitoring Replication Partners

You can monitor information about all the configured replication partners for a group. This information includes both outbound details (volumes on this group replicating to others) and inbound details (replication from other groups to this group).

You should also monitor the usage of delegated space. If free delegated space is not available, replica reserve cannot increase automatically.

Monitoring a Specific Replication Partner

You can monitor details about specific replication partners. For example:

- Partner name, IP address, and contact information
- Status of outbound and inbound replications between a partner and the group
- Details about inbound and outbound replicas and replica collections (for each partner)
- In-process replication activity between groups
- Amount of free delegated space

If free delegated space is low and the replica volume reserve for each replicated volume has not reached its maximum (and, therefore, can increase), consider increasing the delegated space.

About Monitoring Replication Operations

If you are replicating volume data, you should monitor replication operations to make sure that each operation completes.

From the Monitoring tab in the navigation panel, you can see information about:

- Outbound replication (all volumes on the group configured for replication)
- Inbound replication (all replica sets stored in the group from all partners replicating to this group)
- Replication history (history of all outbound replications)
- Replication schedules

In addition, you should monitor the usage of delegated space. If free delegated space is not available, replica reserve cannot increase automatically. You can also monitor replica reserve for a volume. Insufficient replica reserve limits the number of replicas.

Best Practices for Replication Monitoring

Table 64. Replication Monitoring Best Practices identifies some best practices for monitoring replication between groups.

Table 64. Replication Monitoring Best Practices

Monitor	Description
Incomplete replication operations	If a replication operation fails to complete, you might need to increase replication space.
Incomplete manual transfer operations or failback operations	Some operations require multiple tasks that administrators must complete. Make sure you complete all multitask operations.
Low free delegated space	If delegated space is low, replica reserve space might not be able to increase automatically to store new replicas.
Number of replicas	If too many replicas exist, consider decreasing the replica reserve percentage. If too few replicas exist, consider increasing the replica reserve percentage. A low free replica reserve can indicate optimal use of replica reserve space, if the desired number of replicas exist.

Monitor Inbound Replication

Inbound replication transfers volume data from a replication partner to the current group.
To monitor inbound replication:

1. Click **Monitoring**.
2. Select **Inbound Replication**.

The **Inbound Volume Replication** panel displays a list of volumes. Select a volume name to display statistics about that volume.

Monitor Outbound Replication

Outbound replication transfers volume data from the current group to a replication partner.

To monitor outbound replication:

1. Click **Monitoring**.
2. Select **Outbound Replication**.

The **Outbound Volume Replication** panel displays a list of volumes. Select a volume name to display statistics about that volume.

Monitor Outbound Replication History

To see replication history:

1. Click **Monitoring**.
2. Select **Replication History**.

The **Outbound Replication History** panel displays a list of volumes. You should periodically examine the replication duration information. If you see long replication times, make sure that the network connection between the partners is sufficient. A slow network link between the partners can cause long replication times. If a replication operation makes no progress, the group generates a warning event. Make sure that you have adequate network bandwidth between the groups, in addition to full IP routing. If necessary, increase the network bandwidth.

In the **Data size** column of the panel, check how much data you are replicating. You might want to use manual transfer replication if you are transferring a large amount of data.

Monitor Replication Schedules

To monitor replication schedules:

1. Click **Monitoring**.
2. Under Schedules, select **Replication Schedules**.

To see more detail about a schedule, move the pointer over a schedule entry in the panel. A pop-up window appears, showing additional information:

- Partner for the replication
- Schedule type (once, daily, hourly)
- Next date and time the schedule runs
- Schedule status (enabled, disabled, or expired)
- Replication partner name
- Number of replicas to keep

You can take the following actions on a schedule:

- To modify a schedule, either double-click the schedule, or select it and click **Modify**. The **Modify schedule** dialog box opens. Make the changes, then click **OK**. You can change the following items:
  - Schedule name
  - Run date and time
  - Number of replicas to keep
- To disable or enable a schedule, select it, right-click in the highlighted area, and then select the appropriate option from the menu. If you are disabling a schedule, click **Yes** in the confirmation dialog box.
- To delete a schedule, select it and click **Delete**. In the confirmation dialog box, click **Yes** to delete the schedule.

Monitor Replication Partners

You can display information about all the configured replication partners for a group. This information includes both outbound details (volumes on this group replicating to others) and inbound details (replication from other groups to this group).

To display a list of all the replication partners for a group:

- Click **Replication** → **Replication Partners**.

The Replication Partners panel shows all replication partners, replication direction (inbound and outbound), replication status, delegated space, and free space.

You should monitor the usage of delegated space. If free delegated space is not available, replica reserve cannot increase automatically.

Monitor a Specific Partner

To display details about a specific partner:

1. Click **Replication**.
2. Select the replication partner.

The General Partner Information panel shows the partner name, IP address, status, and contact information.

The Volume Replication Status panel shows the status of outbound and inbound replications between this partner and the group. In this panel, check the amount of free delegated space. If free delegated space is low and the replica volume reserve for each replicated volume has not reached its maximum (and, therefore, can increase), consider increasing the delegated space.

For each partner, you can display the following details:

- Outbound replicas
- Outbound replica collections
- Inbound replicas
- Inbound replica collections

Monitor NAS Replication Schedules

1. Click **Monitoring** in the navigation menu.
2. Under NAS Schedules, select **NAS Replication Schedules**.
3. If you need to change a schedule, click one of the icons in the Actions column:
   - Modify
   - Enable (if disabled) or Disable (if enabled)
   - Delete

Monitor Alarms and Operations

The Alarms and Operations panel at the bottom of the GUI displays a visual cue to alarm conditions in the group, as well as in-progress operations. Some operations might need administrator intervention.

To open or close the Alarms and Operations panel, click the panel header or the arrow in the header.

Monitoring Alarms

The group generates an alarm if a persistent hardware condition occurs in a member (for example, high temperature or a failed power supply). Alarms help you discover and correct problems before they disrupt operations. Make sure you investigate all alarms. Each alarm has a priority level, based on the severity of the problem:

- Warning — Condition that decreases performance or can become critical if you do not correct it.
- Critical — Serious problem that can cause damage to the array or data loss.

When an alarm occurs:

- The Alarms panel header flashes. Click the header to open and close the panel.
- The group generates a corresponding event message.
- LEDs on the array chassis light.

The Alarms panel header is divided into two areas: Alarms and Operations. Each header includes icons that match the tabs in the panel.

Alarms header icons are as follows:

- Critical (red circle with an X) and a count of all critical alarms
- Warning (yellow triangle with an exclamation mark) and a count of all warning alarms
- Actions (light bulb) with a count of all actions needed

Operations header icons are as follows:

- Group operations (gear) with a count of all in-process operations
- Failback operations (volume cylinder with an arrow) and a count of all in-process failback operations

Each alarm entry includes the severity, the member that reported the alarm (if applicable), and the message text. Move the pointer over the message text to display more information.

Alarms remain in the Alarms panel until you correct the condition or complete the task. However, the event message associated with the alarm remains in the event log even after the task is complete or the condition is corrected.

Click the **Acknowledge all** icon (✔) in the Alarms panel to acknowledge all alarms.

Display the Alarms Panel

The Alarms and Operations panel at the bottom of the GUI window displays information about alarms in the group and tasks that are in progress. Click the center of the black bar to display the panel.

Table 65. Alarm Panel Icons describes the icons on the Alarms and Operations panel, along with the appropriate user action.

Table 65. Alarm Panel Icons

Icon	Action
⇧ or ⇩	Click the icons or the title bar to open and close the Alarms panel. Each panel contains a link to the object (member or volume) that you can click for additional information.
✔	Click to acknowledge all alarms.
⚠	Flashes to indicate an unacknowledged alarm. Click the acknowledge alarm icon to make this icon stop flashing.

Display Critical Alarms

Open the Alarms and Operations panel and click the **Critical** tab to display Critical alarms. A critical alarm indicates a serious problem that can cause damage to the array or data loss. Correct the problem that causes a critical alarm immediately.

Table 66. Alarm and Operations Panel — Critical Tab shows the contents of the Alarm and Operations panel.

Table 66. Alarm and Operations Panel — Critical Tab

Field	Description
Severity	Severity of the alarm
Object	Object to which the alarm applies
Condition	Condition that triggered the alarm

Critical alarms correspond to ERROR events. These alarms include, but are not limited to:

- Data integrity:
  - RAID is not functioning
  - More than one valid RAID set in the array
  - Full lost block table
- Cache:

  - Control module cache has lost data
  - Cache battery is not charging because it exceeds the temperature limit
  - Cache contains data that does not belong to any of the installed disk drives
- Cooling component fault:
  - Array temperature exceeds upper or lower limit
  - Missing fan tray or cooling module
  - Both fans failed on a fan tray or cooling module
- Hardware component fault:
  - Failed NVRAM coin cell battery
  - Control modules are different models
  - Failed critical hardware component
  - Missing or failed operations panel (not all array models)
  - Failed array monitoring processor (not all array models)

Display Warning Alarms

A warning alarm indicates a condition that decreases performance or can become critical if it is not corrected.

1. Click the up arrow icon at the bottom left of the GUI window.
2. Open the Alarms and Operations panel and click the **Warnings** tab to display Warning alarms.

Table 67. Alarms and Operations Panel – Warning Tab shows the contents of the Alarm and Operations panel.

Table 67. Alarms and Operations Panel – Warning Tab

Field	Description
Severity	Severity of the alarm
Object	Object to which the alarm applies
Condition	Condition that triggered the alarm

Warning alarms correspond to WARNING events. These alarms include, but are not limited to:

- Data integrity:
  - Degraded but functioning RAID set
  - RAID (volume-level) has lost blocks
  - Installed spare drive does not have enough capacity to replace a RAID set drive
- Hardware component:
  - Failed non-critical hardware component
  - Component temperature is near upper or lower limit
  - Fan RPMs exceed upper or lower limit
  - Failed power supply fan
  - Missing power supply
  - Power supply does not have power
- Control module:
  - One installed control module
  - Control module failover occurred
  - Control module has insufficient RAM
  - Lock on secondary control module is open (not all array models)

  - Active control module syncing with secondary
  - No communication between control modules
- Batteries:
  - Real-time-clock battery has low charge
  - Cache battery has less than 72 hours of charge

Monitor Group Operations

1. Open the Alarms and Operations panel.
2. Click the **Group Operations** tab to view the group management operations (for example, moving a member to another pool) and actions that you might need to take.

The Alarms and Operations panel displays details about in-process operations in the group, including moving volumes or members to another pool, moving a partner's delegated space to another pool, or deleting a member. Depending on the operation, you might be able to perform actions on it, such as canceling it.

Monitor Failback Operations

1. Open the Alarms and Operations panel.
2. Click the **Failback Operations** tab to display failback operations and any actions that you might need to take.

Monitor Storage Pool Free Space

You must maintain sufficient free pool space to ensure that load-balancing, thin-provisioning, member-removal, snapshot, and replication operations perform optimally.

To monitor the storage pool free space:

1. Click **Group**.
2. Click **Storage Pools** to open the Storage Pool summary window.

Check the free pool space value in the **Storage Pools** panel. Dell recommends that free pool space not fall below the following level (whichever is smaller):

- 5 percent of the total pool space
- 100GB multiplied by the number of pool members

You can increase free pool space by moving volumes from the low-space pool to a different pool.

You can expand pool capacity by moving a member to the pool. You can sort the table by clicking the heading of a column; the table is sorted by Volume Template by default.

About Diagnostics

Diagnostics collect information about events that occur on a Dell EqualLogic storage system. Dell customer support can use these generated diagnostic reports to propose a solution to a problem with your system.

You can generate a complete diagnostic report based on diagnostics run on PS Series and FS Series controllers. By default, both of these diagnostic types are run on all controllers, enabling your support representative to analyze the most complete diagnostic report for your system. These diagnostic reports are stored on both the FS Series and PS Series arrays. The PS Series diagnostic report can be emailed automatically to your Dell customer support representative. Due to the size of the FS Series diagnostic report, that diagnostic report can be stored only on the array.

- Generate diagnostic reports only when your support representative instructs you to do so.
- Change the default settings only if your support representative instructs you to do so.

You can run reports for the following diagnostics:

- PS Series—Run on individual controllers on each array. These reports are automatically generated and stored on the array and can be emailed to your Dell customer support representative.
- FS Series—Run on all controllers. These NAS cluster diagnostic reports are automatically generated and stored on the array. Because of the size of these reports, they cannot be emailed to your Dell customer support representative, but can be retrieved by File Transfer Protocol (FTP) or Service Control Point (SCP).

If you are troubleshooting a problem with a PS Series system and have SupportAssist enabled, SupportAssist might provide a more detailed diagnostic report on the system. Contact Dell Technical Support to find out if SupportAssist could help you.

The FluidFS Diagnostic Tool provides a quick customer troubleshooting tool to determine any network or other issues. Customers should run this tool only when instructed to by support. After running this tool, customers should send the diagnostic logs to support. See the *Dell EqualLogic Group Manager CLI Reference Guide* for information about the tool.

Add Diagnostic Contact Information

You can configure your system or group to send system diagnostic information to Dell customer support and other email recipients. The group can send PS Series diagnostic information to the specified recipients using multiple email messages.

You will need the following information:

- SMTP IP addresses — SMTP server addresses are used in the listed order if the default server is not available. You can specify up to three addresses. The format for SMTP IP addresses is: 123.45.67.89.
- Sender's email address — This address appears in the From field of the message sent. The format is: *your_name@domain.tld*.

Optionally, you can configure Email Home notification to specify a support provider address to send the diagnostic reports. See Configure or Change Email Home Notifications for more information.

Generate and Email Diagnostic Reports

If an issue arises that only Dell Technical Support can correct, your support representative might instruct you to collect encrypted diagnostic information from one or more group members.

NOTE:

- If you are troubleshooting a problem with a PS Series system and have SupportAssist enabled, SupportAssist might provide a more detailed diagnostic report on the system. Contact Dell Technical Support to find out if SupportAssist could help you.
- If you want to make NAS cluster diagnostics available through FTP, you must enable FTP on the NAS cluster.

To generate diagnostic reports on a system without NAS configured:

1. Expand **Tools** and click **Diagnostic reports**. The Welcome tab opens.
2. Click **Next**. The Getting Started tab opens.
3. If you are working with Dell support and have been assigned a case number, type that number in the **Case Number** field.

   If you have not contacted Dell Technical Support, leave this field blank.
4. Click **Next**.
5. Verify that all members are selected to be included in the diagnostic report. A member is selected if the Select checkbox for that member is checked.

   If your Dell Technical Support representative directed you to remove any member from the report, verify that the Select checkbox for that member is clear.
6. Click **Next**.
7. On the **PS Series Diagnostics – Select Members** tab, select the member to include in the report and click **Next**.
8. On the **PS Series Diagnostics – Report Destination** tab, specify email information for sending the generated report:

   - If you are working with a Dell support representative and want to email a diagnostic report to your representative, select **Send reports to your support provider**.
   - If you need to specify other email addresses where you want to send diagnostic reports, select **Send reports to other email addresses** and enter the email addresses.

- If no SMTP server is configured, click **Configure SMTP** and add the IP address of the SMTP server that you want to configure.

9. Click **Next**. The Summary tab opens.
10. Review the information and then click either:
    - **Copy** to save the summary information in a text file
    - **Back** to return to a previous page to change the diagnostics report settings
    - **Finish** to perform the specified diagnostics and generate the report

To generate diagnostic reports on a system with NAS configured:

1. Expand **Tools** and click **Diagnostic reports**. The Welcome tab opens.
2. Click **Next**. The Getting Started tab opens.
3. If you are working with Dell Technical Support and have been assigned a case number, type that number in the **Case Number** field.

   If you have not contacted Dell Technical Support, leave this field blank.
4. Verify that all types of diagnostics are selected (the default).

   If your Dell support representative directed you to eliminate any type of diagnostics, verify that the checkbox is clear.
5. Click **Next**. If you are generating PS Series diagnostics, the PS Series Diagnostics – Select Members tab opens.
6. Verify that all members are selected to be included in the diagnostic report. A member is selected if the Select checkbox for that member is checked.

   If your Dell support representative directed you to remove any member from the report, verify that the Select checkbox for that member is clear.
7. Click **Next**.
8. If you are generating PS Series diagnostics, on the PS Series Diagnostics – Select Members tab, select the member to include in the report and click **Next**.
9. If you are generating PS Series diagnostics, on the PS Series Diagnostics – Report Destination tab, specify email information for sending the generated report:
    - If you are working with a Dell Technical Supportand want to email a diagnostic report to your support representative, select **Send reports to your support provider**.
    - If you need to specify other email addresses where you want to send diagnostic reports, select **Send reports to other email addresses** and type the email addresses.
    - If no SMTP server is configured, click **Configure SMTP** and add the IP address of the SMTP server that you want to configure.
10. Click **Next**. If you are generating NAS cluster diagnostics, the NAS Cluster – Report Destination tab opens.
11. Click **Next**. The Summary tab opens.

    NOTE: The Summary tab is the only place that IP address and port information for FTP and SCP is available:
    - For PS Series diagnostics, the email address to which the PS Series diagnostic report will be sent and IP addresses and ports where diagnostic reports can be accessed by FTP and SCP.
    - For FS Series diagnostics, the IP addresses and ports where NAS controller diagnostic reports can be accessed by FTP and SCP.
12. Review the information and then click either:
    - **Copy** to save the summary information in a text file
    - **Back** to return to a previous page to change the diagnostics report settings
    - **Finish** to perform the specified diagnostics and generate the report

To view the status of the diagnostics on each member, expand Operations at the bottom of the window, and click the **Group Operations** tab.

Enable FTP on a NAS Cluster

If you are working with a Dell customer support representative to run diagnostics on NAS controllers, you must enable FTP to provide your representative access to the NAS diagnostics report.

To enable FTP on a NAS cluster:

1. Click **Group**, expand **Group Configuration**, and select the NAS cluster.
2. Click the **Advanced** tab.
3. In the NAS Cluster Access panel, select the **Enable FTP Access** checkbox.

Troubleshooting Performance Issues

To effectively manage your storage performance, follow these basic steps to troubleshoot issues:

- Fix any failed hardware components in the Dell storage array:
  - Ensure that the array is fully populated (drives, controllers, and power supplies).
  - Replace any failed disks or failed control modules.
- Configure all network interfaces:
  - Make sure all the network interfaces on the group members are configured and functional.
  - Configure data center bridging (DCB) on your network to make use of iSCSI.
- Fix or upgrade network hardware:
  - Correct any network hardware problems that are causing poor performance.
  - Fix switches or add network bandwidth, if needed.
  - Add new hardware for better performance.
- Fix or upgrade server hardware:
  - Correct any server hardware problems that are causing poor performance.
  - Add memory, update drivers or network adapters, or use multipathing I/O, if needed to improve performance.
- Optimize applications:
  - Correct any application behavior that is causing poor performance.

    The goal is to reduce the load on the group. For example, be sure database volumes are provisioned correctly. You might want to place data and log files on different volumes.
- Optimize your PS Series group configuration:
  - Verify that your pools have sufficient capacity by adding members.
  - Verify that your RAID policy is appropriate for a group member. For more information about RAID policies on PS Series systems, review the Dell Technical Report entitled *PS Series Storage Arrays: Choosing a Member RAID Policy*, which you can download from the following location: http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19861480.
  - Use higher-performing disk drives.
  - Ensure that all group members have the latest PS Series firmware.

Server and Application Performance Recommendations

You might be able to improve performance by following these general recommendations:

- Make sure you have sufficient server memory. Adding memory can decrease the read I/O load on the storage.
- Make sure you have the latest drivers for your operating system.
- Add high-performance network interfaces or host bus adapters to servers.
- Consider using multipathing (or MPIO), which provides a high-performance and highly available path between a server and storage. Use the Dell EqualLogic Host Integration Tools for a multipathing solution.
- Make sure the server Ethernet ports, PCI bus, and CPU are rated for the workload.
- If the server is a clustered computer, isolate cluster network traffic from iSCSI storage traffic. Check for other network traffic interference.
- Make sure the application is configured properly so that I/O is performed efficiently.

Network Infrastructure Performance Recommendations

Network performance is complex and depends on a number of components working with each other. You might be able to improve network performance by following these general recommendations:

- Make sure that network components are recommended for an iSCSI SAN by Dell EqualLogic.
- Make sure the switches and interswitch links have sufficient bandwidth for the iSCSI I/O. Contact EqualLogic customer support for details on correct sizing and configuration of iSCSI SAN switch infrastructure. Pay careful attention to sizing interswitch links.
- Make sure all member network interface connections are Gigabit Ethernet and make sure driver settings are correct on the servers. (SAN Headquarters generates an alert if it detects connections that are less than 1GB.)
- Make sure you connect and enable all the member Ethernet interfaces to maximize the available SAN bandwidth. If all interfaces are enabled but bandwidth is still insufficient, increasing the number of arrays in the storage pool might provide additional throughput if the servers have not reached their maximum bandwidth.
- If only one interface is completely utilized on a member or on a server with multiple NICs, ensure that MPIO is properly configured. If all of the server NICs exceed their capacity (to determine this, use host-based tools such as SAN Headquarters) but the PS Series group has excess network capacity, add additional server NICs. Also, configure MPIO for those operating systems that support MPIO.
- Lack of receive buffers can also cause network performance problems. Low-end switches often have limited memory and suffer from performance issues related to insufficient buffers. Dell recommends a 1MB per port buffer level in switches and, preferably, dedicated buffers over shared buffers.
- Server performance can often be improved by increasing the number of buffers allocated to the server NICs. Consult your switch vendor or server NIC driver documentation to determine if you can increase the buffers.
- Network bandwidth saturation is typically not a problem if it occurs during sequential operations and if application performance remains acceptable. To resolve problems with network bandwidth, you might have to redistribute the workload over multiple arrays. The PS Series architecture permits bandwidth to scale, either by enabling additional array interface ports or by adding another array to the group and thus adding more controllers and network ports.
- Make sure you also follow the network performance guidelines in the *Installation and Setup Guide*. In particular:
  - If possible, do not use Spanning Tree Protocol (STP) functionality on switch ports that connect end nodes.
  - Enable flow control on switches and NICs — The most critical function that affects network performance is flow control, which allows network devices to signal the next device that the data stream should be reduced to prevent dropped packets and retransmissions.
  - Disable unicast storm control on switches.
  - Enable jumbo frames with care — In environments with small average I/O sizes, using jumbo frames has limited benefits. In general, support for jumbo frames is disabled by default on switches and server NICs. Enabling jumbo frames requires that the switch use a VLAN other than the default VLAN (usually VLAN 1). PS Series arrays automatically negotiate using jumbo frames when the iSCSI connection is established by the server.

    Consult your switch vendor or server NIC driver documentation to determine if jumbo frames can be configured. Note that some network devices run more slowly with jumbo frames enabled, do not properly support jumbo frames, or cannot support them simultaneously with flow control. In these cases, jumbo frames should be disabled, or the switches should be upgraded or replaced.

PS Series Group Performance Recommendations

To improve performance on your PS Series groups, consider the following group configuration recommendations for these areas: network interfaces, storage pools, RAID policy, volumes, hardware and firmware, iSCSI connections, and MPIO.

Network Interfaces

Make sure all the network interfaces on the members are configured, functioning, and accessible. Use the Group Manager CLI ping command to check accessibility.

NOTE: If you need more network bandwidth for iSCSI I/O, do not configure a dedicated management network.

Storage Pool Capacity

Low storage pool capacity is a problem that generates an alert in SAN Headquarters. If a pool has less than 5 percent of free space (or less than 100GB per member, whichever is less), a PS Series group might not have sufficient free space to efficiently perform the virtualization functions required for automatic optimization of the SAN. In addition, when storage pool free space is low, write performance on thin- provisioned volumes is automatically reduced to slow the consumption of free space.

If pool capacity is low, try one or more of the following remedies:

- Move volumes from the low-space pool to a different pool.
- Reduce the amount of in-use storage space by deleting unused volumes or by reducing the amount of snapshot reserve.
- Increase pool member capacity by fully populating the drive bays or upgrading to higher-capacity disks.
- Increase pool capacity by adding a member.
- Move delegated space to a different pool.

RAID Policy

Consider changing the RAID policy for a member. Change the policy only if you are sure your applications will perform better with a different RAID level. RAID 10 performs better than other RAID levels when a disk drive fails and when a RAID set is degraded. In addition, RAID policies with spare disks provide additional protection against drive failures.

NOTE: You can change a RAID policy for a member only if the new RAID policy uses less disk space than the current RAID policy.

Volume Management

- Assign application volumes to a pool that includes members with the RAID policy that is optimal for the application.
- Automatic load balancing — Consider not binding volumes to a particular group member or using fewer pools, and let the group perform automatic performance load balancing.

Hardware and Firmware

- Replace member disk drives with higher-performing drives. Make sure the array is fully populated with disk drives.
- Make sure member control module caches are in write-back mode.
- Make sure all the group members are running the latest PS Series firmware.

iSCSI Connections

Large, complex environments can use many iSCSI connections. A storage pool in a PS Series group can support numerous simultaneous connections, as described in the *Dell EqualLogic PS Series Storage Arrays Release Notes*. These connections can be used for fully provisioned volumes, thin-provisioned volumes, and snapshots.

Attempting to exceed the supported number of connections results in an error message. You can reduce the number of iSCSI connections to the volumes and snapshots in a storage pool in several ways:

- Disconnect from unused volumes and snapshots.
- Modify MPIO settings to reduce the number of connections per volume.
- Move volumes to another storage pool.
- Create a new storage pool and move volumes to the new storage pool.

Multipath I/O

MPIO provides additional performance capabilities and network path failover between servers and volumes. For certain operating systems (Windows 2003 and 2008), the connections can be automatically managed. If MPIO is not creating multiple connections, you should:

- Verify that the storage pool does not have the maximum number of iSCSI connections for the release in use. See the *Dell EqualLogic PS Series Storage Arrays Release Notes*.
- Verify the access control policies for the volume. Using the iSCSI initiator name instead of an IP address can make access controls easier to manage and more secure.
- Ensure that Dell EqualLogic MPIO extensions are properly installed on the supported operating systems. See the Dell EqualLogic Host Integration Tools documentation for details.

- Ensure that MPIO is supported and properly configured, according to the documentation for the operating system.

You can also monitor multiple PS Series groups with SAN Headquarters and can launch the Group Manager GUI from there; however, you cannot directly manage the storage from SAN Headquarters.

A

Third-Party Copyrights

All third-party copyrights for software used in the product are listed below.

This product contains portions of the NetBSD operating system:

For the most part, the software constituting the NetBSD operating system is not in the public domain; its authors retain their copyright.

Copyright © 1999-2001 The NetBSD Foundation, Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

All advertising materials mentioning features or use of this software must display the following acknowledgement:This product includes software developed by the NetBSD Foundation, Inc. and its contributors.

Neither the name of the NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS “AS IS”' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This code is derived from software contributed to The NetBSD Foundation by Charles M. Hannum and by Jason R. Thorpe of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

This code is derived from software contributed to The NetBSD Foundation by John T. Kohl and Charles M. Hannum.

This code is derived from software contributed to The NetBSD Foundation by Kevin M. Lahey of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

This code is derived from software contributed to The NetBSD Foundation by Jun-ichiro Hagino.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

Copyright © 1995, 1996, 1997, 1998 Christopher G. Demetriou.

This code is derived from software contributed to The NetBSD Foundation by Luke Mewburn.

This code is derived from software contributed to The NetBSD Foundation by Klaus Klein.

This code is derived from software contributed to The NetBSD Foundation by Jonathan Stone.

This code is derived from software contributed to The NetBSD Foundation by Jason R. Thorpe.

This code is derived from software contributed to The NetBSD Foundation by UCHIYAMA Yasushi.

This product includes software developed for the NetBSD Project by Wasabi Systems, Inc.

Copyright © 2000-2001 Wasabi Systems, Inc. All rights reserved.

This product includes software developed by the University of California, Berkeley and its contributors. This product includes software developed by the University of California, Lawrence Berkeley Laboratory.

Copyright 1985-1995 The Regents of the University of California.

Copyright 1997-2000 Niels Provos.

This code is derived from software contributed to Berkeley by Ralph Campbell.

This code is derived from software contributed to Berkeley by Rick Macklem.

Copyright © 1989 Digital Equipment Corporation.

This product includes software developed by Manuel Bouyer.

Copyright © 1999 Manuel Bouyer.

This product includes software developed by Adam Glass.

Copyright © 1994 Adam Glass.

This code is derived from software contributed to Berkeley by Paul Vixie.

This code is derived from software contributed to Berkeley by Chris Torek.

This code is derived from software contributed to Berkeley by Mike Hibler.

This code is derived from software contributed to Berkeley by Paul Borman at Krystal Technologies.

This code is derived from software contributed to Berkeley by Peter McIlroy.

This code is derived from software contributed to Berkeley by Peter McIlroy and by Dan Bernstein at New York University.

This code is derived from software contributed to Berkeley by Stephen Deering of Stanford University.

This code is derived from software contributed to Berkeley by Jeffrey Mogul.

Copyright 1996 The Board of Trustees of The Leland Stanford Junior University.

This product includes software developed by the Computer Systems Laboratory at the University of Utah. Copyright © 1990,1994 The University of Utah and the Computer Systems Laboratory (CSL). All rights reserved.

This code is derived from software contributed to Berkeley by the Systems Programming Group of the University of Utah Computer Science Department.

Copyright (c) 2000 Soren S. Jorvang.

Copyright (c) 1993 John Brezak. All rights reserved.

Copyright © 1995 - 2000 WIDE Project. All rights reserved.

© UNIX System Laboratories, Inc.

All or some portions of this file are derived from material licensed to the University of California by American Telephone and Telegraph Co. or UNIX System Laboratories, Inc. and are reproduced herein with the permission of UNIX System Laboratories, Inc.

Copyright © 1999 Shuichiro URATA.

This product includes software developed by Matthias Pfaller.

Copyright © 1996 Matthias Pfaller.

Copyright © 1993 Jan-Simon Pendry.

This product includes software developed by Gordon W. Ross.

Copyright © 1995 Gordon W. Ross.

This product includes software developed by Philip A. Nelson.

Copyright © 1993 Philip A. Nelson.

Copyright © 1999 Ross Harvey.

This product includes software developed by Christos Zoulas.

Copyright © 1996 Christos Zoulas.

Copyright © 1997 Zubin D. Dittia.

This product includes software developed by SiByte, Inc.

Copyright © 2000 SiByte, Inc.

Copyright © 1996, 2000 Intel Corporation.

Copyright 1996 - 1998 Microsoft Corporation.

Copyright © 1990,1994 The University of Utah and the Computer Systems Laboratory (CSL).

Copyright © 1991 Bell Communications Research, Inc. (Bellcore).

Copyright © 2000 Caldera Systems, Inc. All rights reserved.

Copyright © 1995 - 2000 Kungliga Tekniska Högskolan.

(Royal Institute of Technology, Stockholm, Sweden). All rights reserved.

Copyright © 1993-1995 HEWLETT-PACKARD COMPANY

Copyright © 1995-1997 Eric Young. All rights reserved.

Copyright © 1992 Simmule Turner and Rich Salz. All rights reserved.

Copyright © 1999 - 2001, PADL Software Pty Ltd. All rights reserved.

Copyright © 1985 - 1988 by the Massachusetts Institute of Technology.

Copyright © 1995 by Wietse Venema. All rights reserved.

Copyright © 1999 The OpenSSL Project. All rights reserved.

Copyright © 1992 – 1999 Theo de Raadt. All rights reserved.

Copyright © 1999 Dug Song. All rights reserved.

Copyright © 2000-2002 Markus Friedl. All rights reserved.

Copyright © 2001 Per Allansson. All rights reserved.

Copyright © 1998 CORE SDI S.A., Buenos Aires, Argentina.

Copyright © 2001-2002 Damien Miller. All rights reserved.

Copyright © 2001 Kevin Steves. All rights reserved.

Copyright © 1999 Aaron Campbell. All rights reserved.

Copyright © 2002 Nils Nordman. All rights reserved.

Copyright © 2000 Todd C. Miller. All rights reserved.

Copyright © 1995, 1996 by David Mazieres.

Copyright © 2000 Zembu Labs, Inc. All rights reserved.

Copyright © 2000 Takuya SHIOZAKI. All rights reserved.

Copyright © 1992 Keith Muller.

Copyright © 1994, Jason Downs. All rights reserved.

Copyright © 1997 Matthew R. Green. All rights reserved.

Copyright © 1999 Citrus Project. All rights reserved.

Copyright © 1990-2, RSA Data Security, Inc. All rights reserved.

Copyright © 1995 by International Business Machines, Inc.

Copyright © 1996 by Internet Software Consortium.

Copyright © 1995, 1999 Berkeley Software Design, Inc. All rights reserved.

Copyright © 1993 Carlos Leandro and Rui Salgueiro Dep. Matematica Universidade de Coimbra, Portugal, Europe.

Copyright © 1992, 1993, 1994 Henry Spencer.

Copyright © 1986-1991 by Sun Microsystems, Inc.

Copyright © 1993 Martin Birgmeier.

Copyright © 1991 by AT&T.

Copyright © 1997 Frank van der Linden. All rights reserved.

Copyright © 1999 Michael Graff. All rights reserved.

This product includes software developed by Alistair G. Crooks.

Copyright © 1999 Alistair G. Crooks. All rights reserved.

Copyright © 2001 Cerbero Associates Inc.

Copyright © 1995-1998 Mark Adler.

Copyright © 1995-1998 Jean-loup Gailly.

Copyright © 1998-1999 Brett Lymn. All rights reserved.

Copyright © 1996-1999 SciTech Software, Inc.

Copyright © 2001,2002 Brian Stafford.

Copyright © 1999-2001 Bruno Haible.

Copyright © 2001 Alex Rozin, Optical Access. All rights reserved.

Copyright © 1989 TGV, Incorporated. All rights reserved.

Copyright © 2000 Frank Strauss. All rights reserved.

Copyright © 1997 Niels Baggesen. All rights reserved.

Copyright © 2000 National Aeronautics & Space Administration. All rights reserved.

Copyright © 1990-2000 Robin’s Nest Software, Inc.

Copyright © 1989-1996 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by Carnegie Mellon University. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.