EQUALLOGIC

PS Series Storage Arrays

Group Administration

PS Series Firmware Version 5.0

June 2010

Part Number: 110-6027-EN_R1

Table of Contents

Preface

This document describes PS Series group functionality and management operations. It describes how you use the Group Manager graphical user interface (GUI) to perform tasks.

For information about using the command line interface (CLI) to manage a group, see the *CLI Reference* manual on the EqualLogic customer support Website.

Audience

This documentation is for administrators responsible for managing a PS Series group or SAN storage and also for individuals interested to learn more about group operation. You can manage a PS Series array without extensive network or storage experience. It is important that understand:

- Basic networking concepts – IP addresses, network bandwidth, and network interfaces.
- Current network environment – How your LAN is currently configured such as whether it is organized into subnets, whether you use an NTP server, and the type of iSCSI initiators that you use.
- User disk storage requirements – How you intend to assign storage space on the group to your user community and the kind of service level agreement you maintain for data recovery.
- RAID configurations – The method of disk-level data security that you want to implement through a RAID policy
- Disk storage management – Principles for implementing data security and recovery.

Note: Although this documentation includes examples of PS Series arrays in some common network configurations, information about setting up a network is beyond the scope of this manual. Go to the EqualLogic customer support Website to find more information about network requirements.

Organization

Part I: General Group Management:

- Chapter 1, *Storage solutions for all enterprises*, includes an introduction to PS Series arrays, group features, and functionality.
- Chapter 2, *Common group tasks*, describes some common post-setup tasks that Dell recommends.
- Chapter 3, *Group Manager user interfaces*, describes the user interfaces for managing a group and how to access online help.
- Chapter 4, *Group security*, describes how to use accounts to protect your group from unauthorized access and how to set up a dedicated management network.
- Chapter 5, *Group configuration*, describes how to modify the basic group configuration.
- Chapter 6, *Group members*, describes how to set the RAID policy for a member and configure the member network configuration.

Part II: Using Group Storage:

- Chapter 7, *Storage pools*, describes how to organize storage for a group.
- Chapter 8, *iSCSI target security*, describes how to protect volumes and snapshots from unauthorized and uncoordinated iSCSI initiator access.
- Chapter 9, *Basic volume operations*, describes how to create volumes and perform basic volume tasks.
- Chapter 10, *Advanced volume operations*, describes thin provisioning, template and thin clone volumes, and advanced volume management tasks.
- Chapter 11, *Snapshot management*, describes how to create and manage volume snapshots.
- Chapter 12, *Volume replication*, describes how to create replicas and manage replication.
- Chapter 13, *Data recovery*, describes how to recover data from replicas.

Part III: Troubleshooting

- Chapter 14, *Group event logging*, describes how the group logs events.
- Chapter 15, *Group monitoring*, describes how to monitor events, interpret status, identify and solve problems, and monitor performance statistics.

In addition, this guide includes a *Glossary* of terms used in PS Series group management and related product offerings from EqualLogic.

Conventions

Documentation conventions are shown in the following table.

Document Conventions

Convention	Usage
`fixed width`	Command, parameter, output, file name, link, button, field, URL address, or e-mail address.
group_ip_address	Variable. Replace the text in italics with the actual object name or identifier.
Note: *Text . . .*	Important information that you should read.
Requirement: *Text . . .*	Information relating to one or more requirements for performing a task.
Recommendation: *Text . . .*	Information relating to one or more Dell recommendations for performing a task.
Restriction: *Text . . .*	Information relating to one or more restrictions for performing a task.

Overview of EqualLogic Products

Thank you for your interest in EqualLogic™ PS Series storage products from Dell. We hope you will find them intuitive and simple to configure and manage.

PS Series arrays optimize resources by automating performance and network load balancing. Additionally, PS Series arrays offer all-inclusive array management software, host software, and free firmware updates. The features and products described next are available at no additional cost.

PS Series Array Software

- Firmware – Installed on each array, PS Series firmware allows you to manage your storage environment and provides capabilities such as volume snapshots, cloning, and replication to protect data hosted on the array in the event of an error or disaster.
  - Group Manager GUI: Provides a graphical user interface for managing a group.
  - Group Manager CLI: Provides a command line interface for managing a group.
- Manual Transfer Utility (MTU) – Runs on Windows and Linux systems and enables you to use physical media to securely transfer large amounts of data to a replication partner, facilitating replication and preventing network congestion.

Host Software for Windows

- Host Integration Tools:
  - Remote Setup Wizard (RSW): Initializes new PS Series arrays, configures host connections to a group, and configures and manages multipathing.
  - Multipath I/O Device Specific Module (MPIO DSM): Includes a connection awareness-module that understands PS Series network load balancing and facilitates host connections to PS Series volumes.
  - VSS and VDS Provider Services: Allows backup software vendors to perform off-host backups.
  - Auto-Snapshot Manager/Microsoft Edition (ASM/ME): Uses PS Series snapshots, cloning, and replication to provide point-in-time protection of critical data for supported applications, including SQL Server, Exchange Server, Hyper-V, and NTFS file shares.
- SAN HeadQuarters (SANHQ): Provides centralized monitoring, performance trending, and event reporting for multiple PS Series groups.

Host Software for VMware

- Storage Replication Adapter for Site Recovery Manager (SRM): Allows SRM to understand and recognize PS Series replication for full SRM integration.
- Auto-Snapshot Manager/VMware Edition (ASM/VE): Integrates with VMware Virtual Center and PS Series snapshots to enable Smart Copy protection of Virtual Center folders, datastores, and virtual machines.

Current Customers: You may not be running the latest versions of the software listed above. If you are under a valid warranty or support agreement for your PS Series array, you are entitled to obtain the latest updates and new releases as they become available.

Related Documentation

For detailed information about PS Series arrays, groups, volumes, array software, and host software, see the documentation listed in the following figure:

Figure 2-1: PS Series Documentation

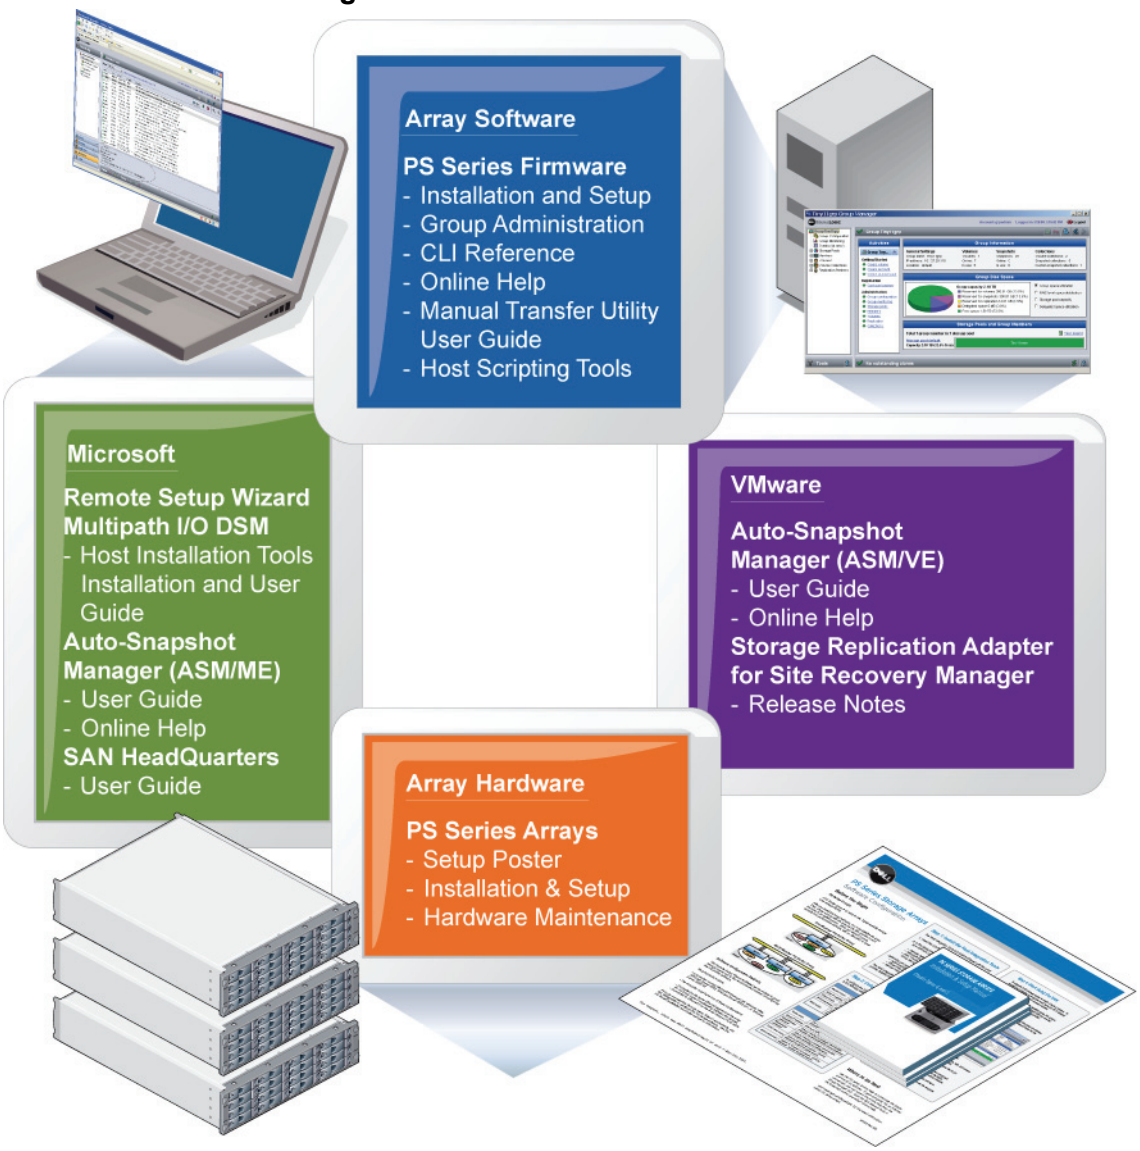

Technical Support and Customer Service

Dell's support service is available to answer your questions about PS Series arrays. If you have an Express Service Code, have it ready when you call. The code helps Dell's automated-support telephone system direct your call more efficiently.

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.

For customers in the United States, call 800-945-3355.

Note: If you do not have access to an Internet connection, contact information is printed on your invoice, packing slip, bill, or Dell product catalog.

Use the following procedure to contact Dell for sales, technical support, or customer service issues:

1. Visit support.dell.com or the Dell support URL specified in information provided with the Dell product.
2. Select your locale. Use the locale menu or click on the link that specifies your country or region.
3. Select the required service. Click the "Contact Us" link, or select the Dell support service from the list of services provided.
4. Choose your preferred method of contacting Dell support, such as e-mail or telephone

Online Services

You can learn about Dell products and services using the following procedure:

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

1 Storage solutions for all enterprises

PS Series storage arrays provide consolidated storage in a self-managing, iSCSI (Internet Small Computer System Interface) storage area network (SAN). Featuring automated management and fast, flexible scalability, PS Series arrays can greatly decrease the total cost of storage acquisition and management.

Today, large and small businesses alike are under pressure to manage rapidly-growing storage environments. Some still use storage directly attached to servers, which can be a a difficult configuration to scale that offers low service levels. The EqualLogic storage solution provides the benefits of storage consolidation to enterprises of all sizes and types.

To use storage resources efficiently and economically, you can configure disparate storage configurations so administration is easier and service levels increase. The EqualLogic storage solution also provides centralized data provisioning, backups, and disaster recovery features to simplify SAN management.

About PS Series groups

The foundation of the EqualLogic storage solution is the **PS Series group**, which includes one or more PS Series arrays (**group members**) you can connect to an IP network and manage as a single system. Each array has fully redundant hardware and multiple network connections for high availability and I/O bandwidth.

You can connect multiple PS Series arrays as a group to an IP network and manage them as a single storage system, while leveraging your current network infrastructure. With PS Series arrays, you can deploy a full-featured iSCSI SAN. Access to storage requires only an iSCSI initiator, which is available with most operating systems.

Virtualization technology masks the underlying complexity of the storage configuration, saving you time and effort. Arrays share configuration data and “collaborate” to achieve automatic data provisioning and load balancing.

Self-managing features of PS Series arrays enable straightforward configurations that do not disrupt running applications. For example, you can add more arrays to a group to seamlessly increase capacity and performance. You can automatically replicate data over long distances for a simple, yet robust, disaster recovery strategy.

Integrated virtualization software provides:

- Group Manager configuration, management, and replication software
- Automatic RAID and spare disk drive configuration
- Automatic network, performance, and capacity load balancing

How groups work

You might start with a single-member group and later add arrays to expand capacity and increase performance. Whether the group is small or large, there is no disruption to users, and management overhead is static. The group automatically distributes data across member disk drives and network interfaces. Load-balancing, based on capacity and performance metrics, occurs as needed.

Administration accounts provide security for group management operations. You can set up a dedicated management network for additional security, if required.

By default, a group provides a single pool of storage. If you have multiple members, you can divide group space into different **storage pools** and then assign members. Pools help you organize storage according to usage and give you a way to allocate resources, from a single-system management view.

You create **volumes** to access storage space in a pool, and you can modify volume size and attributes on-demand. **Cloning** a volume enables you to make an exact copy of the volume.

Figure 1-1 shows a three-member group (three network connections on each member), with members and volumes distributed across two pools, and two servers, with iSCSI initiators accessing the volumes.

Figure 1-1: Three-Member, Two Pool PS Series Group

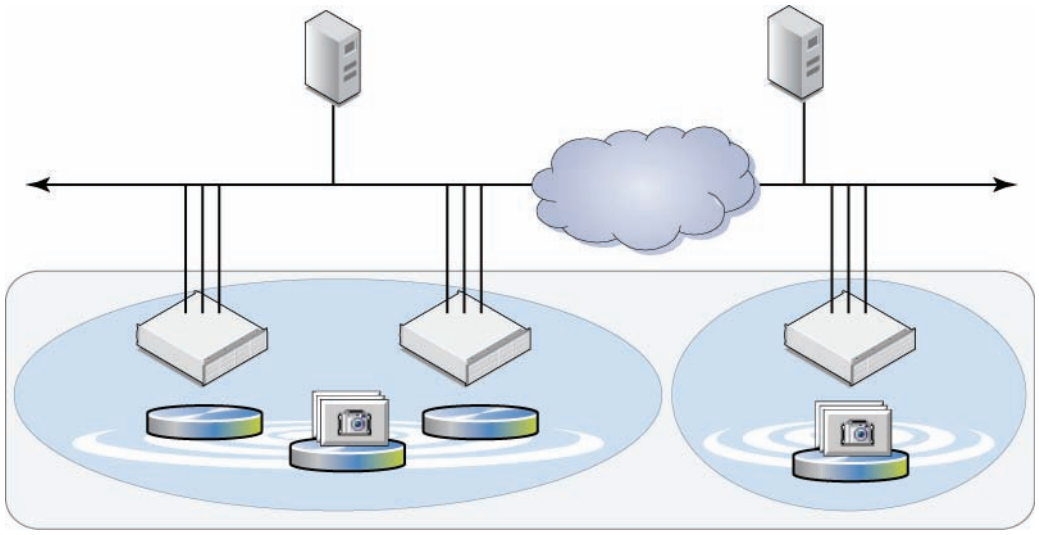

High-end features in an affordable iSCSI san

To meet the storage requirements of businesses today, a PS Series group offers advanced storage capabilities in a highly-available, cost-effective iSCSI SAN.

Modular hardware

A PS Series group can easily adjust to workload requirements; therefore, administrators purchase hardware only when necessary. New product versions fully operate with previous arrays.

Depending on the model, a PS Series storage array provides the following features:

- **Redundant hardware for no single point of failure.** Redundant, hot-swappable hardware components—disks, control modules, fans, and power supplies—offer the highest availability. Component failover and disk sparing occur automatically without user intervention or disrupting data availability. In addition, data in each array is kept safe by RAID technology.
- **High-performance control modules.** Each control module has multiple Gigabit Ethernet interfaces. You can configure multiple network interfaces for automatic failover. With dual control modules, the array mirrors data between battery-backed write-back caches. If one control module fails, the other control module starts operating—automatically and with no disruption to users.

- **Support for standard Ethernet networks.** Only one IP network connection is necessary for array operation. You can configure all the network interfaces for maximum bandwidth. You do not have to train administrators in unfamiliar and complex technologies such as Fibre Channel. Also, high-volume production and intense vendor competition among Ethernet hardware vendors decrease cost of ownership.
- **Easy and online expansion.** You can expand SAN capacity with no disruption to users and applications. Old array models can work in groups with new models, ensuring the viability of your initial purchase.
- **Optional management connectivity through a serial port.** If network connectivity is lost, a serial connection—providing full management capabilities—lets you connect the array to a console, terminal server, or computer running a terminal emulator and manage the storage.
- **SAS, SATA, and SSD disk drives.** A variety of disk drive speeds and capabilities enable you to optimize your storage environment to meet your business needs.

Seamless online scalability

To increase the capacity of an individual PS Series array, you can install more drives or configure more network connections. Expand overall group capacity by adding arrays to the group. In all cases, performance scales linearly, and the new disk and network resources are immediately available for use. The additional control modules also increase processing power. Meanwhile, volumes stay available with no affect on computers and applications, and management overhead stays the same.

Interoperability

PS Series arrays are ideal for multi-platform, heterogeneous environments that previously required a different storage system for each operating system or application. Because a PS Series group provides block-level storage, computers can use a standards-compliant iSCSI initiator—available for most major operating systems—to access data. When a computer connects to a volume, it looks like a regular disk that you can format using the normal operating system utilities.

Easy setup and management

You can quickly configure an array on the network and create a PS Series group. In minutes, you have a functioning iSCSI SAN. Automation of complex operations like RAID configuration, disk sparing, data provisioning, and load balancing means that even novices can effectively manage the SAN.

You can configure arrays and groups by using a serial connection to an array and running the `setup` utility. Alternately, Windows computers can run the Remote Setup Wizard to configure an array and set up access to the SAN. See *Host-based applications* on page 1-6 for more information.

Password-protected management accounts give easy, secure access to the group. Graphical and command-line user interfaces give you a single-system view of the storage. With the Group Manager graphical user interface (GUI), creating and managing volumes, configuring security, and setting up event notification are simple and intuitive operations. You can access the same functionality by using the Group Manager command-line interface (CLI) through telnet, SSH, or a serial connection.

Automatic SAN operation

In contrast to traditional storage management environments involving error-prone, manual tasks, a PS Series group does complex tasks correctly and without user intervention:

- **Automatic RAID configuration and data provisioning.** Administrators do not have to manually create RAID sets or map data onto disks or individual network interfaces. Whether you are expanding a group or creating, expanding, or deleting volumes, the group manages storage allocation and load balancing across the resources in the group. iSCSI access to volumes continues without interruption.
- **Automatic spare disk configuration and use.** A PS Series array can include spare disks, which the array automatically configures and uses in the event of a disk failure, offering "hands-off" storage management.
- **Automatic event notification.** A PS Series group uses standard event logging mechanisms and can automatically notify users of significant events through e-mail, remote syslog servers, or SNMP traps.

Dynamic load balancing

Dynamic load balancing lets the group quickly find and correct bottlenecks as the workload changes, with no user intervention or application disruption.

A group provides three types of load balancing within the arrays in a storage pool:

- **Capacity load balancing.** The group distributes volume data across disks and members, based on capacity.
- **Performance load balancing.** The group tries to store volume data on members with a RAID configuration that is optimal for volume performance, based on internal group performance metrics.
- **Network connection load balancing.** The group distributes iSCSI I/O across network interfaces, minimizing I/O contention and maximizing bandwidth.

Robust security for group administration

The web-based Group Manager GUI and its corresponding CLI enable you to manage a group. You can access the GUI by using standard Web protocols. For increased security, you can use SSL encrypted Web access. You can access the CLI by using telnet. For increased security, you can use SSH.

Password-protected administration accounts prevent unauthorized users from accessing the group. The `grpadmin` account is the default administration account. You can create additional accounts with a variety of privileges.

You can set up and authenticate administration accounts locally in the group or through a RADIUS (Remote Authorization Dial-in User Service) server. Environments that want additional security can use a dedicated management network, which enables you to separate management traffic from iSCSI traffic.

Robust security for data access

Access controls prevent unauthorized users from accessing volume data. The robust security mechanisms of the iSCSI protocol mean that you do not have to understand complicated security technologies, such as Fibre Channel Switch Zoning or LUN Masking.

You can restrict iSCSI access to volumes according to IP address, iSCSI initiator name, or Challenge Handshake Authentication Protocol (CHAP) user name. You can set up and authenticate CHAP accounts locally in the group or through a RADIUS (Remote Authorization Dial-in User Service) server.

Advanced functionality at no extra cost

PS Series arrays match ease-of-use with ease-of-doing-business. Arrays come with basic and advanced software, so you do not have to purchase expensive add-ons. New features are available through firmware updates and host-based tools, which are provided at no additional cost to customers with a support contract.

For example, a PS Series group delivers:

- **Online scalability.** You can expand array capacity or overall group capacity while online and with no affect on users or availability. Hot-swappable hardware means that you can replace failed components while data stays online.
- **Multiple storage pools.** In a PS Series group, you can divide storage into multiple pools. This helps you organize storage according to usage, offering more control over resource allocation, while giving you a single system management view. You get the advantages of storage consolidation and the ability to easily segregate different workloads.
- **Thin provisioning.** To use storage more efficiently, you can apply thin provisioning to a volume. Thin provisioning lets you over-allocate group space to satisfy current and future storage requirements. This functionality can be useful if an operating system or application cannot handle online volume expansion.

  To an application or operating system, a thin-provisioned volume is fully allocated. However, the group initially allocates only a portion of the volume size. As you use the volume, the group automatically allocates more space, with no disruption in availability.
- **Thin clones.** To use storage space efficiently, you can create multiple thin clone volumes that are based on a template volume. You can write data to each thin clone to make it unique. Only the data that is unique to a thin clone is consumed from free pool space. Thin clones can be beneficial in storage environments that use multiple volumes that contain a large amount of common data.
- **Volume collections**. You can organize volumes into collections. This lets you perform a snapshot or replication operation on all the volumes in the collection simultaneously.
- **Scheduled operations.** You can set up schedules to create volume snapshots or replicas at a specific time in the future or on a regular basis.
- **Cloning.** Cloning a volume creates a new volume with the same size, contents, and attributes as the original volume. Cloning is useful when you need a copy of an existing volume, such as when you are deploying copies of a computer or database. When compared to traditional copy or restore operations, cloning dramatically decreases the time required to make complete copies of a volume. You can also clone a snapshot or a replica, creating a new volume with the same contents as the volume at the time you created the snapshot or replica.
- **Snapshots.** A snapshot is a point-in-time copy of volume data. Snapshots greatly simplify and increase the performance of backup and recovery operations. You can create snapshots on-demand and through schedules.

  To recover data from a snapshot, you can set the snapshot online, connect to the target, and copy the data. You can also restore the entire volume from a snapshot or clone a snapshot. Cloning a snapshot creates a volume containing the same data that was in the volume at the time the snapshot was created.
- **Replication.** By replicating volumes from one group to another, you can set up a simple, yet robust disaster recovery strategy. Groups can be in the same building or a large distance apart.

A replica represents the contents of a volume at a specific point in time. You can create replicas on-demand and through schedules.

- **Failover and failback.** 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 failback to the original group.

Host-based applications

Host-based applications facilitate access to PS Series group storage and expand group capabilities at no extra cost. These applications include user interfaces that provide interaction between the computer and the group. See *Overview of EqualLogic Products* on page xiv.

Part I: Managing Groups

2 Common group tasks

When you use the setup utility or the Remote Setup Wizard to create a new PS Series group with one or more members, you have a fully functioning group with many features.

Dell recommends you perform the common post-setup tasks listed here. See *About the Group Manager GUI* on page 3-1 for information about logging in to the group.

Setting the group 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, EST (Eastern Standard Time).

Group time is based on the first member's clock. Each array you add updates its clock to match the group's date and time. The group date and time determines the timestamp which is used to identify some objects (snapshots, for example) you create on the group.

You can change the group time manually (see *Modifying the time zone and clock time* on page 5-3).

You can also specify that the group use an external Network Time Protocol (NTP) server to automatically set the same time for all members (see *Setting the time through an NTP server* on page 5-4).

Creating 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. Dell recommends one account per administrator. See *Creating a local administration account* on page 4-5 for more information.

Setting up event notifications

Dell recommends that you configure event notification, so you automatically receive messages when events occur in the group. Events enable you to track operations and also detect and solve problems before they affect performance or data availability.

See *Group event logging* on page 14-1 and select one or more notification methods to automatically receive notification when events occur:

- **E-mail notification.** If an event occurs, the group automatically sends a message to designated e-mail addresses. See *Configuring E-Mail notification* on page 14-3.
- **E-Mail Home.** If a hardware component fails or if you update firmware, the group automatically notifies customer support. E-Mail Home is available to all PS Series customers, but response time and assistance is based on the validity and level of your support contract. See *Configuring E-Mail home* on page 14-4.
- **Remote server logging.** The group logs events to a remote server at the syslog facility. You can also access events from the syslog server. See *Configuring syslog notification* on page 14-5.

Configuring CHAP for initiator authentication

You can use Challenge Handshake Authentication Protocol (CHAP) for iSCSI authentication to manage access controls more efficiently. Using a challenge-response mechanism, CHAP restricts target access through user names and passwords instead of unique IP addresses or iSCSI initiator names. You can use CHAP to authenticate iSCSI initiators by specifying a CHAP user name in an access control record. To meet this condition, a computer must supply the user name and its password (or "secret") in the initiator configuration interface when logging in to the target.

See Chapter 8, *iSCSI target security*.

Configuring SNMP access to the group

You can use Simple Network Management Protocol (SNMP) for read-only group access. In addition, some monitoring tools (such as SAN HeadQuarters) and the Manual Transfer Utility require SNMP access to the group.

See *Displaying and configuring SNMP access to a Group* on page 4-11.

Setting group-wide volume defaults

When you create a volume (or enable thin-provisioning for a volume), the group applies default values to volume settings that control snapshot space, snapshot behavior, thin provisioning space, and iSCSI alias naming.

You can change the default values to meet the needs of your environment.

See *Displaying group-wide default volume settings* on page 9-5.

Setting the RAID policy for a member

To use the storage in a group member, you must set the member RAID policy.

See *Setting the RAID policy and pool for a new member* on page 6-6.

Configuring member network interfaces

After you add a member to a group, the member has one configured network interface, typically Ethernet 0. For best performance and availability, Dell recommends that you configure all the network interfaces that are eligible for iSCSI traffic.

For network configuration requirements and recommendations, see the *Hardware Maintenance* manual for your array model.

To configure member network interfaces, see *Configuring network interfaces* on page 6-9.

3 Group Manager user interfaces

PS Series arrays provide simple, yet robust user interfaces for creating, expanding, and managing groups. The graphical user interface (GUI) and the command-line user interface provide virtually identical functionality.

About the Group Manager GUI

You can run the Group Manager GUI:

- From a Web browser. In addition to ports 3002 and 3003, the GUI uses the standard HTTP port (80).
- By installing the GUI on a local computer and running it as a standalone application. See *Installing and starting the GUI application* on page 3-1.

See the PS Series *Release Notes* for the latest information on GUI requirements.

Starting the GUI from a Web browser

1. Enter the group IP address (or the dedicated management network address) in a supported Web browser.
2. At the login prompt, enter a valid group administration account name (for example, the `grpadmin` account) and password.

In addition, you can install the Group Manager GUI on a computer and run the GUI as a standalone application. See *Installing and starting the GUI application* on page 3-1.

Installing and starting the GUI application

You can install the Group Manager GUI on a computer and run the GUI as a standalone application. You can install GUIs for more than one group on a computer. You identify GUI applications by group name.

Requirement: Make sure the computer on which you installed the GUI application is running the required Java version. See the *PS Series Storage Arrays Release Notes* for details.

Note: See the PS Series *Release Notes* for the latest information on standalone GUI requirements.

1. Using the computer where you installed the GUI, launch the Group Manager GUI. See *Starting the GUI from a Web browser* on page 3-1.
2. Expand `Tools` in the lower-left portion of the GUI and click `Run as application`.
3. Confirm 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.
4. At the login prompt, enter a valid group administration account name and password.
5. Start the standalone GUI application, using one of the following methods (This sequence might vary, depending on your operating system):

   a. Click the shortcut associated with the group name.

b. Navigate to the Windows `Programs` menu, and click `EqualLogic PS Group,` then *`group_name`* `Group Manager.`

When you install the GUI locally, it automatically updates when you upgrade 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 configure a management network for the group, you must uninstall the GUI application and then install it again.

Uninstalling the GUI application

1. From the Windows Control Panel, click `Java`.
2. In the `General` tab, under `Temporary Internet Files`, click `View`.
3. Right-click the GUI application for the group and select `Delete`. The name uses the format: *`group_name`* `Group Manager`.

   Depending on the Java version, the GUI application for a group might have two components--the applet and the library; both are prefaced with the group name and show EqualLogic as the vendor. Make sure you remove both components.

Navigating the GUI

The first time you log in to the GUI, the Group Summary window shows an overview of the group configuration and status. Any alarms are shown in the panel at the bottom of the window. See Figure 5-1 on page 5-1.

Note: The PS Series firmware GUI can also show replication storage objects that originate from other PS Series groups (replication partners).

The GUI provides a set of view buttons in the bottom left of the window:

- Group view (default) - Provides a view of the group, organized as:
  - Group information and group configuration
  - Storage pools
  - Members (physical arrays that are part of this group)
- Volumes - Provides a view of storage objects, organized as:
  - Volumes - a hierarchy of storage objects such as standard volumes, template volumes and thin clones and their associated snapshots.
  - Volume collections - user-specified related sets of storage objects on which you can perform operations such as creating snapshots.You can also configure a regular schedule to perform operations automatically.
  - Custom snapshot collections - user-specified related sets of snapshots (independent of volume collections).
- Replication - Provides a view of the disaster recovery replication configuration for the group, organized as:
  - Replication partners - other PS Series groups that receive replicated data from this group or are authorized to send replicated data to the login group. (The term *login group* describes a group that you are currently logged into.)

  - Inbound replica collections - related sets of storage objects located on a partner group, sending replicated data to the login group.
  - Inbound replicas- individual storage objects located on a partner group, sending replicated data to the login group.
  - Out bound replica collections - related sets of storage objects located on the login group, sending replicated data to a partner group.
  - Outbound replicas - individual storage objects located on the login group, sending replicated data to a partner group.
- Monitoring - Provides a view of group information, status, and messages, organized as:
  - Events - The group event log messages, filtered by log event class.
  - Statistics - Events and data relating to administrative login sessions and iSCSI initiator connections.
  - Schedules - Events and data relating to snapshot and replication schedules.
  - Replication - The status of storage objects configured for replication on the login group and on replication partner groups. This view also provides a cumulative log showing all replication events.

Keyboard shortcuts

Table 3-1 lists keyboard shortcuts you can use to navigate the GUI without using a mouse. The first column lists the pane in the window and the second pane lists the keyboard shortcut for it.

Note: There are also keyboard shortcuts for individual buttons and fields in the GUI. See the GUI online help for more information.

Table 3-1: Keyboard Shortcuts

Type	Pane	Shortcut
General	Switch to Group view	Control+Alt+G
	Switch to Volumes view	Control+Alt+V
	Switch to Replication view	Control+Alt+R
	Switch to Monitoring view	Control+Alt+M
	Toggle the Tools panel	Control+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	Control+S
	Discard all changes	Control+Z
	Refresh data	Control+R
	Move to next item	Tab
	Move to previous item	Shift+Tab
	Open the Help Context menu	F1

Table 3-1: Keyboard Shortcuts (Continued)

Type	Pane	Shortcut
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	Control+Tab
	Leave table and move to the previous item in the pane	Shit+Control+Tab
	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	Control+Page Up
	Next Tab	Control+Page Down
Alarms	Show/hide Alarms panel	Control+Alt+A
	Acknowledge All button	Control+Shift+K

GUI icons

Table 3-2 identifies the icons at the top of the GUI window. The first column lists the icons and the second describes them.

Table 3-2: GUI Icons

Icon	Description	Keyboard Shortcut
	**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.	Control+S
	**Discard changes**. Discards changes you made in a window.	Control+Z
	**Refresh data in a window**. Refreshes the data that appears in the GUI. Do not use the browser refresh button to refresh the data that appears in the GUI.	Control+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)

Accessing the alarms panel

The Alarms panel at the bottom of the GUI window shows alarms in the group and tasks that are in progress.

- Click the Show window icon () or Hide window icon () in the Alarms title bar or click the title bar to open and close the panel. Each alarm contains a link to the object (member or volume). Click the link for additional information.
- Click the Acknowledge all icon () to acknowledge all alarms and stop flashing the Caution icon ().

Displaying the tools panel

If the Tools panel is not showing in the bottom left of the GUI window, click its Show window icon ( ) to open the Tools panel. To keep the Tools window open, drag the panel divider bar .

Table 3-3 shows the Tools panel configuration options and utilities for working in the Group Manager:

Table 3-3: Tools Panel

Option	Description	User Actions
User preferences	Opens the Modify user preferences dialog box that enables you to: • Set General GUI Policies • Set General GUI Policies • Set Alarm Policies • Set Data Validation and Debugging Policies	See *Customizing the GUI* on page 3-5
Online help	Opens a new browser window displaying topic-oriented help on every aspect of using the GUI to manage a group.	See *Starting online help for group manager* on page 3-7
Customer support	Launches the EqualLogic Customer Support Website.	See *Contacting customer support* on page 15-33
Performance monitor	Opens the performance monitor.	See *Using the Performance Monitor* on page 15-27
Manual transfer utility	Opens the Manual Transfer Utility (MTU).	See *About manual transfer replication* on page 12-3
Run as application	Installs the Group Manager GUI on the local computer, enabling it to run as a standalone application. (If you are currently running the GUI as an application, this link does not appear.)	If the GUI is already installed, clicking the link starts the application. See *Installing and starting the GUI application* on page 3-1
Diagnostic reports	Generates Diagnostic reports that are sent out via e-mail to the address configured for e-mail notification.	See *Collecting diagnostic information* on page 15-33
Update firmware	Opens the Update Firmware dialog box.	See *About member firmware* on page 6-12

Customizing the GUI

The topics that follow describe how you can:

- Change the look and behavior of the GUI, and the location of online help.
- Specify what the GUI does when a connection is lost.
- Change event logging options and how alarms are indicated.
- Control the GUI response to typed input and debugging.

Setting general GUI policies

To set general GUI policies to control the appearance of the GUI and the location of online help:

1. Click `Tools`, then `User Preferences`, then the `General` tab.

2. In the User Preferences – General dialog box, select your preferences and click `OK`.

See the online help for information about the options.

Setting GUI communication policies

To manage connections between your workstation and the Group Manager GUI:

1. Click `Tools`, then `User Preferences`, then the `Communication` tab.
2. Select the policies and click `OK`.

See the online help for information about the options.

Setting alarm policies

To set an alarm which will indicate there is a hardware condition in an array:

1. Click `Tools`, then `User Preferences`, then the `Alarms` tab.
2. In the User Preferences – Alarms dialog box, select policies and click `OK`.

See the online help for information about the options.

Setting advanced policies

To set data validation and debugging policies:

1. Click `Tools`, then `User Preferences`, then the `Advanced` tab.
2. In the User Preferences – Advanced dialog box, select policies and click `OK`.

Using the CLI

The Group Manager command-line interface (CLI) provides a comprehensive set of commands for managing a PS Series group. The CLI also enables you to manage individual PS Series storage arrays for maintenance purposes.

To access the group to run CLI commands, you can do one of the following:

- Use a 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.
- Use a 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 (the ACT LED is green). Use the serial cable that shipped with the array. See the *Hardware Maintenance* manual shipped with your array for more information.

See the *CLI Reference* manual for more information about using the CLI. You can download documentation from the EqualLogic Customer Support Website.

Starting online help for group manager

In addition to tooltips and command-line help for the 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 Web site in the Dell.com domain. You have the option to install the help on your local system or a private Web server.

Note: Setting the online help to a local path only applies to users of the local system.

The online help is context-sensitive. Click any of the question mark icons embedded in the Group Manager GUI to open a specific help topic. You can also open the help in a browser, and use the table of contents or search option to find information.

To launch online help from the Web site, in the GUI's far-left panel, expand `Tools` and click `Online Help`.

To launch online help locally (from your system):

1. Log on to the Customer Support Web site and navigate to the downloads area.
2. Copy the folder named `eqlgmhlp`*`nn`* (where *`nn`* corresponds to the PS Series firmware release, such as 5.0).
3. Save the contents of folder to on a local disk, a network share, or Web server repository. The top-level file in the help hierarchy in this folder is named `groupmanager.htm`. For example: `c:/eqlgmhlp50`.
4. Start the Group Manager GUI.
5. Click `Tools`, then `User Preferences`, then `General tab`.
6. In the User Preferences – General dialog box, enter the new help directory location in the `Location of online help files field` and click `OK`. For example:
   ```
   file:///c:/eqlgmhlp50
   http://servername/eqlgmhlp50
   file:///system.directory.company.com/myhelpshare/eqlgmhlp50
   ```
7. Click any help icon to test the revised help location.

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. Start Internet Explorer and go to `Tools`, then `Internet Options`.
2. Select the `Security` tab and click `Trusted Sites`, then `Sites`.
3. Add the group's IP or management address to the list of trusted sites, using the format: `http://`*`group_ip_address`*. (Do not select the option to require server verification.)
4. Close both dialog boxes.

4 Group security

Group security features enable you to control access to the group and the data it contains.

About group security

To access a group for management purposes, an administrator must meet several security conditions. See Table 4-1.

Table 4-1: Access Requirements for Group Administration

Security Condition	Description
Network access	The administrator's computer must have access to the group network address (group IP address or dedicated management address).
Group administration access enabled in the group	To use the GUI, the group must allow administrative access through the web. To use the CLI, the group must allow administrative access through telnet or SSH.
Valid group administration account	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.

In addition to administration account security, Table 4-2 identifies other group security options.

Table 4-2: Group Security Options

Security Option	Description
RADIUS authentication	You can control access to a group and its volumes by using administration accounts to log in to the group. Using a RADIUS Authentication server enables you to centralize account management.
SNMP	Simple Network Management Protocol (SNMP) enables read-only access to the group.
VDS/VSS access control	Enables Windows VDS and VSS access to the group. You must create at least one VDS/VSS access control record that matches the access control credentials you configure on the computer by using Remote Setup Wizard or Auto-Snapshot Manager/Microsoft Edition.
Dedicated management network	An advanced option enables you to configure a dedicated management network, which separates group management traffic from iSCSI traffic.

Accessing the GUI or CLI

By default, administrators can access the GUI remotely using a Web browser or a standalone Java application. Administrators can also manage a group by using the command-line interface (CLI) across a telnet or SSH connection.

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

Note: 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 re-enable access. See *Using the CLI* on page 3-6 and your *Hardware Maintenance* manual for information about serial connections.

Administration access options

Table 4-3 shows the access options and network services.

Table 4-3: Administration Access Options

Field	Description	Shortcut	User Actions
Enable Web access	Whether administrators can access the Group Manager through the web interface. **Note:** To run Auto-Snapshot Manager/ VMware Edition, you must enable use of the GUI for group administration.	Alt+W	See *Enabling or disabling GUI or CLI access* on page 4-2
Allow only secure SSL connections	Whether administrators must use an SSL connection when connecting to the GUI.	Alt+A	See *Enabling or disabling GUI or CLI access* on page 4-2
Enable telnet access	Whether the CLI can be accessed through a telnet connection.	Alt+T	See *Enabling or disabling GUI or CLI access* on page 4-2
Enable SSH (secure shell) access	Whether the CLI can be accessed through an SSH connection.	Alt+S	See *Enabling or disabling GUI or CLI access* on page 4-2

The CLI provides additional options for managing network services. Using the CLI, you can also enable or disable SSH V1 protocol support or the ftp service. See the *CLI Reference* manual.

Enabling or disabling GUI or CLI access

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.
2. In the Administration Access panel, enable or disable the GUI or CLI access options and network services as described in *Administration access options* on page 4-2.
3. Click `Save all changes` (Control+S).

About administration accounts

Having persons in the role of administrator is important to protect and maintain your group from unauthorized access. Environments that need additional security might also benefit from a dedicated management network.

To manage or monitor a group, you must log in to an administration account. Administration accounts prevent unauthorized individuals from accessing a group.

An account can authorize an individual to perform all group operations, perform only operations on a pool (and optionally monitor the entire group), manage its own volumes within an assigned quota, or only monitor the group, depending on the type of account.

The default administration account, `grpadmin`, can perform all group operations. Dell recommends that you set up an account for each administrator.

Recommendation: Dell recommends one account per user and that the group administrator monitors the activity of other accounts. See Chapter 15, *Group monitoring*, for more information.

You can manage accounts locally or remotely:

- Locally in the group – If you have relatively few administration accounts, this method is practical. Account authentication occurs within the group. The default administration account, `grpadmin`, is a local account created automatically when the group is first configured.

  See *Creating a local administration account* on page 4-5.

- Remotely on an external 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 and, optionally, manage administration accounts.

  Restriction: To delete a RADIUS account, remove it from Active Directory and then delete it from the group.

  A group can use both local accounts and RADIUS-authenticated accounts. However, each account name must be unique.

  See *About administration accounts on a RADIUS authentication server* on page 4-7.

Types of administrator accounts

Table 4-4 lists administration account types and their privileges, The first column lists account types and the second column describes them.

Table 4-4: Types of Administrator Accounts

Account Type	Description
`grpadmin`	Can perform all group management tasks, including managing the group, storage pools, members, volumes, and accounts. You set the password for the `grpadmin` account when you create a group. You cannot delete the `grpadmin` account. Only the `grpadmin` account can update member firmware. 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 cannot update member firmware.
Read-only	Can view information about all group objects, but cannot change the group configuration.
Pool administrator	Can manage the volumes, members, snapshots, and other objects only in the pool or pools for which the account has authorization. Optionally, pool administrators can view information about all group objects. Pool administrators can assign volumes to volume administrators, provided that the pool administrator has access to the pool containing the volumes, and that the volume administrator has sufficient free quota space. Pool administrators cannot change the resources to which they have access.

Table 4-4: Types of Administrator Accounts (Continued)

Account Type	Description
Volume administrator	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 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. 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. 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, volumes, or replication partners to which they have access.

Administrators accounts have these restrictions:

- You cannot change the name of an administration account. Instead, you must delete the account and then re-create it with the new name.
- You cannot disable, delete, change the name, or change the type of the `grpadmin` account.
- Only group administrator accounts can change the attributes of accounts, with the exception of the `grpadmin` account restrictions above.
- Volume administrator, pool administrator, and read-only accounts can only change the password, description, and contact information for their accounts.

Administration account attributes

Table 4-5 displays the attributes of administration accounts. The first column lists the attributes, the second column describes them. Gather this information before creating an account.

Table 4-5: Administration Account Attributes

Attribute	Description
Name	Name of the account, up to 16 alphanumeric characters, including period (.), hyphen (-), and underscore (_). The first character must be a letter or number. The last character cannot be a period.
Password	Password for the account. The password must be from 3 to 16 alphanumeric characters and is case-sensitive. However, validation occurs only for the first 8 characters.
Description	Optional description for the account.

Table 4-5: Administration Account Attributes (Continued)

Attribute	Description
Type	Account type: • Group administrator – Can change any and all aspects of the group, storage pools, members, and volumes, except updating member firmware. • Pool administrator – Can manage the volumes, members, snapshots, and other objects only in the pool or pools for which the account has authorization. Optionally, pool administrators can view information about all group objects. • Volume administrator – Can manage the volumes for which the account has authorization. Additionally, volume administrators can view information about pools to which the account has access. • Read-only – Can view information about all group objects, but cannot change the group.
Managed pools	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 pool(s). Applies to Pool administrators and Volume administrators.
Replication Partners	The group(s) on which the account can delegate space for replication and replicate volumes. Applies to Volume administrators only.
Additional access permission	Grants read access to the entire group. Applies to Pool administrator and Read-only accounts; Volume administrators only have read access to the individual pools containing the storage quota they manage.
Contact	Name, e-mail address, and phone numbers for the account owner.
Enable administration account	Whether the account is enabled or disabled. A user cannot log into a disabled account.

Displaying local administration accounts

To see the names, types, access permissions or status of local administration accounts:

1. Click `Group`, then `Group Configuration.`
2. Select the `Administration` tab. The Group Administration window appears.

See the online help for information about the data fields and options.

Creating a local administration account

You can 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.

Before creating a local administration account, gather the information described in *Administration account attributes* on page 4-4.

1. Click `Group Configuration`, then `Administration tab.`
2. In the Administration Accounts panel, click `Add`. The Create Account - General Settings dialog box appears.
3. Enter the account name, password, and description (optional) and click `Next`. The Create Account - Account Permissions dialog box appears.
4. Select the type of account and (if applicable) the pool access and read access to the group.

- For a pool administrator, select one or more pools the account can manage and whether the account has read-only access to the entire group.
- For a volume administrator, select one or more pools the account can manage and specify the quotas for each pool.

5. Select whether to enable (default) or disable the account, then click `Next`. (You can enable and disable accounts at any time.)

   If you created a volume administrator account, and the group has replication partner(s) configured, the Create Account - Allowed replication partners dialog box appears.

6. [Optional] Select one or more replication partners that this account can replicate to, then click `Next`. the Create Account - Contact information dialog box appears.

7. [Optional] Enter contact information for the account and click `Next`. The Create Account - Summary dialog box opens.

8. Review the account information. Click `Back` to make changes, or click `Finish` to create the account.

Modifying a local administration account

You can modify the account attributes described in Table 4-5. However, you cannot change the account name. Instead, you must delete the account and then re-create it with a new name.

In addition, you cannot disable, delete, change the name, or change the type of the `grpadmin` default administration account.

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.

2. In the Administration Accounts panel, select the account and click `Modify`.

   - To change the account password or description, click the `General` tab and change the information in the Modify Administration Account – General dialog box.
   - To change the account type or pool or volume administrator settings, click the `Permissions` tab and change the information.
   - To change replication partners for a volume administrator, click the `Replication Partners` tab and change the selection(s).
   - To change the account contact information, click the `Contact` tab and change the information.

3. Click `OK`.

Deleting a local administration account

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.

2. In the Administration Accounts panel, select the account and click `Delete`.

3. Confirm that you want to delete the account.

   Note: When you delete a Volume administrator account, the volumes it manages are not deleted, and its replication and operations continue as scheduled.

About administration accounts on a RADIUS authentication server

You can use an external RADIUS authentication server to centralize the management of administration accounts. 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.

Using a RADIUS server can simplify account management if you have a large number of accounts.

There are various implementations of RADIUS, including Microsoft Windows Internet Authentication Service (IAS). Depending on the 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 IAS and 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 web site.

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, or read-only accounts and to store account contact information. See *Types of administrator accounts* on page 4-3 and *Administration account attributes* on page 4-4.

Recommendation: 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:

- `Administrative` – Specifies that the account is either a group administrator account, a pool administrator account, or a volume administrator account.

  Note: If you do not specify the `EQL-Admin` attribute, by default, the account is a group administrator account.

- `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` attribute and the `EQL-Pool-Access` attribute.

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

Table 4-6 describes the Dell vendor-specific attributes for RADIUS attributes, and lists their possible values.

Table 4-6: 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`	12740
	`VSA number`	6
	`VSA syntax`	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` attribute to 0 and the `EQL-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` attribute is 1 (pool administrator account) or 3 (volume administrator account). The quota for volume administration accounts is expressed as *`PoolName Quota`*, with gb and mb appended to the quota representing GB and MB, respectively. 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, e.g. `Pool1 unlimited`. If no unit is specified, the default capacity unit is MB.	`VSA vendor ID`	12740
	`VSA number`	7
	`VSA syntax`	String (comma-separated list of pools; 3 to 247 characters)
`Admin-Repl-Site-Access` Specifies the sites to which the volume administrator can replicate volumes. Required if the value of the EQL-Admin attribute is 3 (volume administrator account). Used only for volume administrators.	`VSA vendor ID`	12740
	`VSA number`	8
	`VSA syntax`	String (comma-separated list of sites; 3 to 249 characters)
`Admin-Account-Type` Specifies whether the account is read-only (`RO`) or read-write (`RW`):	`VSA vendor ID`	12740
	`VSA number`	9
	`VSA syntax`	RO or RW
`Admin-Full-Name` (Optional) Name of the administrator using the account.	`VSA vendor ID`	12740
	`VSA number`	1
	`VSA syntax`	String (3 to 247 characters)

Table 4-6: Vendor-Specific Attributes (Continued)

Attribute	Field	Required Value
`Admin-Email` (Optional) E-mail address of the administrator.	`VSA vendor ID`	12740
	`VSA number`	2
	`VSA syntax`	String (3 to 247 characters)
`Admin-Phone` (Optional) Phone number for the administrator.	`VSA vendor ID`	12740
	`VSA number`	3
	`VSA syntax`	String (3 to 247 characters)
`Admin-Mobile` (Optional) Mobile phone number for the administrator.	`VSA vendor ID`	12740
	`VSA number`	4
	`VSA syntax`	String (3 to 247 characters)
`Admin-Poll-Interval` How often, in seconds, the GUI polls the group configuration data. The default is 30 (seconds).	`VSA vendor ID`	12740
	`VSA number`	5
	`VSA syntax`	Integer (up to 6 numerals)

Displaying RADIUS authentication and accounting servers

1. Click `Group`, then `Group Configuration`.
2. Select the `Administration` tab.

See the online help for information about the data fields and options.

Using RADIUS authentication and accounting servers

Prerequisite tasks for RADIUS servers

Perform the tasks in Table 4-7, in the order listed, before you use a RADIUS server to authenticate administration accounts (or CHAP accounts for iSCSI access, as described in *Using CHAP accounts on a RADIUS authentication server* on page 8-3).

Table 4-7: RADIUS Server Prerequisites

Task	Description
Install and configure the RADIUS authentication server.	For example, to add the group as a RADIUS client on a Network Policy Server, you must specify: • The name (also called *Friendly Name*) for the client. Dell recommends using the group name. • The group IP address, (also called Client address), or dedicated management network IP address. • The Vendor Name attribute. Select RADIUS Standard. • An optional password (also called Shared Secret), of up to 63 characters. This password should also be entered in the Group Manager when you configure the group to use the RADIUS authentication server. Dell recommends that you use a password for increased security. The RADIUS server must be accessible to all the group members.

Table 4-7: RADIUS Server Prerequisites (Continued)

Task	Description
Configure iSCSI CHAP accounts.	For iSCSI CHAP 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 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.
Set up attributes for administration accounts.	For administration accounts, set up the attributes that allow the server to authorize accounts as group administrator, pool administrator, or read-only accounts. See *RADIUS attributes for administration accounts* on page 4-7.
Set up accounts.	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.

Procedure for configuring RADIUS servers

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.
2. Click `RADIUS settings`. The RADIUS Settings dialog box appears.
3. Under RADIUS authentication servers click `Add` and specify the IP address of the server.
4. Specify and then confirm the RADIUS secret.
5. [OPTIONAL] Specify 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 is two seconds.
   - `Number of retries` – Number of times the group tries to contact an accounting server after the first failure. The default is one.
6. [OPTIONAL] Using the Radius accounting servers panel, repeat steps 2 through 5 if you also want to add information for RADIUS accounting servers.
7. Repeat steps 3 through 6 to add additional servers or click `OK` to finish.
8. In the Group Administration window, click `Save all changes` (Control+S).

All RADIUS authentication and accounting options are now enabled. For a description of these options, see *Displaying RADIUS authentication and accounting servers* on page 4-9.

Modifying RADIUS server settings

You can change these settings for a RADIUS authentication or accounting server:

- Server IP address
- Password (secret)
- Request timeout value

- Number of retries value

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.
2. In the RADIUS Authentication panel, click `RADIUS settings`.
3. To change a server IP address or password, select the server IP address and click `Modify` in the RADIUS settings dialog box. Change the settings and click `OK`.
4. Click `OK`.

Deleting a RADIUS server connection

1. Click `Group`, then `Group Configuration`, and then the `Administration` tab.
2. In the RADIUS Authentication panel, click `RADIUS settings`. The RADIUS settings dialog box appears.
3. Select the server IP address and click `Delete`.
4. In the Group Administration window, click `Save all changes` (Control+S).

Disabling use of a RADIUS server in a Group

If you previously configured the use of a RADIUS authentication server or a RADIUS accounting server, you can disable use of the server.

1. In the Group Administration window, deselect `Enable RADIUS authentication for login` or deselect `Enable RADIUS accounting for authenticated users`.
2. Click `Save all changes` (Control+S).

Displaying and configuring SNMP access to a 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 requires you to configure SNMP access and specify `public` for the SNMP community name.

To display SNMP Access:

1. Click `Group`, then `Group Configuration`.
2. Select the `SNMP` tab.

See the online help for information about the data fields and options.

To change SNMP access:

1. In the SNMP Access panel, click `Add`.
2. Enter a SNMP community name (for example, `public`) and click `OK`. You can specify up to five names.
3. Click `Save all changes` (Control+S).

To modify or delete an SNMP community name:

1. In the SNMP Access panel, select the name.
2. Click `Modify` or `Delete`.

Host-based application access requirements

Host-based applications from EqualLogic facilitate access between computers and PS Series group storage.

Each application requires a method of accessing the group and its storage:

- SAN HeadQuarters – Enables you to monitor the performance and status of multiple groups from a central user interface. The application uses SNMP to access the group.

  See the SAN HeadQuarters documentation and *Displaying and configuring SNMP access to a Group* on page 4-11.

- Manual Transfer Utility – Enables you to perform a replication operation using external media, instead of the network. The application uses SNMP to access the group.

  See the Manual Transfer Utility documentation and *Displaying and configuring SNMP access to a Group* on page 4-11.

- Auto-Snapshot Manager/Microsoft Edition – Enables you to create consistent backups on group storage. The application uses special access controls to access the group.

  See the ASM/ME documentation and *Displaying and configuring Windows service access to a Group* on page 4-12.

- Auto-Snapshot Manager/VMware Edition – Enables you to create backups on group storage in a VMware environment. The application uses the Group Manager GUI channel to access the group.

  See the ASM/VE documentation and *Accessing the GUI or CLI* on page 4-1.

Displaying and configuring Windows service access to a Group

Microsoft Windows computers that are running a Microsoft service, such as Virtual Disk Service (VDS) or Volume Shadow Copy Service (VSS), must be able to access a PS Series group to perform management operations.

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

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

VDS/VSS access control records use the same criteria for restricting access as iSCSI target access control records: CHAP user name, iSCSI initiator name, or iSCSI initiator IP address. See *About iSCSI target access controls* on page 8-1.

To display VDS/VSS Access, click `Group`, then `Group Configuration`, and then the `VDS/VSS` tab.

See the online help for information about the data fields and options.

Adding a VDS/VSS access control record

1. Click `Group`, then `Group Configuration`, and then the `VDS/VSS` tab.
2. Click `Add` in the VDS/VSS Access Control List panel.
3. Do at least one of the following:
   - Check the box marked `Authenticate using CHAP user name` to use CHAP (*Authenticating initiators through CHAP* on page 8-2).
   - Check the box marked `Limit access by IP Address` to constrain access to an IP address or range of addresses. Use an asterisk as a wildcard to specify a range of addresses, such as `127.200.*.*`.
   - Check the box marked `Authenticate using CHAP user Limit access to iSCSI Initiator name` to grant access to a specific SCSI initiator (See *iSCSI target security* on page 8-1). For example: `iqn.2000-05.com.qlogic.qla-4000.sn00044.`
4. Click `OK`.

Modifying or deleting a VDS/VSS access control record

1. Click `Group`, then `Group Configuration`, and then the `VDS/VSS` tab.
2. To modify a record, select the record and click `Modify`. Change the CHAP user name, IP address (or range), or the iSCSI initiator name. (See *Adding a VDS/VSS access control record* on page 4-13.) Then, click `OK`.

   To delete a record, select the record and click `Delete`. Then, confirm that you want to delete the record.

When you delete or modify a record you might need to update any computer that was previously accessing volumes through the access control record.

About dedicated management networks (advanced)

For increased security, you can configure a dedicated management network 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, continues to use the group IP address.

Although a dedicated management network can provide additional group administration security, it has disadvantages:

- Because you assign the highest-numbered network interface on each group member to the management network, iSCSI traffic is limited to the remaining network interfaces. Therefore, total iSCSI bandwidth might decrease, depending on the control module type.

  Note: Some control modules have a network interface that you can use only if you configure a dedicated management network. For these control modules, using a dedicated management network does not decrease iSCSI bandwidth. See the Hardware Maintenance manual for your array model.

- 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.

Only very knowledgeable users should configure a management network and only if the environment requires separating management traffic from iSCSI traffic.

Configuring a management network

Before configuring an management network, read the considerations described in *About dedicated management networks (advanced)* on page 4-13.

Prerequisite tasks for configuring 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 is the address 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 that 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.

  See *Configuring a management network* on page 4-14.

- 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`.

Procedure for configuring a management network

Warning: 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.

1. Click `Group`, then `Group Configuration`, and then the `Advanced` tab.
2. Click `Configure management network` in the Dedicated Management Network panel (Alt+C). The Configure Management Network window appears.
3. Select `Enable dedicated management network`.
4. Enter the management network IP address in the `Management IP address` field.
5. Enter the default gateway in the `Default gateway` field.
6. For each group member:

   a. Double-click to configure, and enable at least one network interface, other than the highest-numbered interface, on the iSCSI network.

   For the best performance connect, configure, and enable all iSCSI-eligible interfaces. See *Configuring a management network* on page 4-14.

   b. Click `Configure for management-only access` next to the highest-numbered network interface.

   c. In the Modify IP Settings – Management Network dialog box:

   – Enter an IP address that is on the management network subnet.
   – Enter a subnet mask (netmask).
   – Select `Enable this interface`.
   – Select `Restrict to management access`.
   – Click `OK`.
7. Verify the network configuration in the Configure Management Network dialog box (Figure 4-1).

   Click `OK` to complete the dedicated management network configuration.

Figure 4-1: Configure Management Network – Configuration Complete

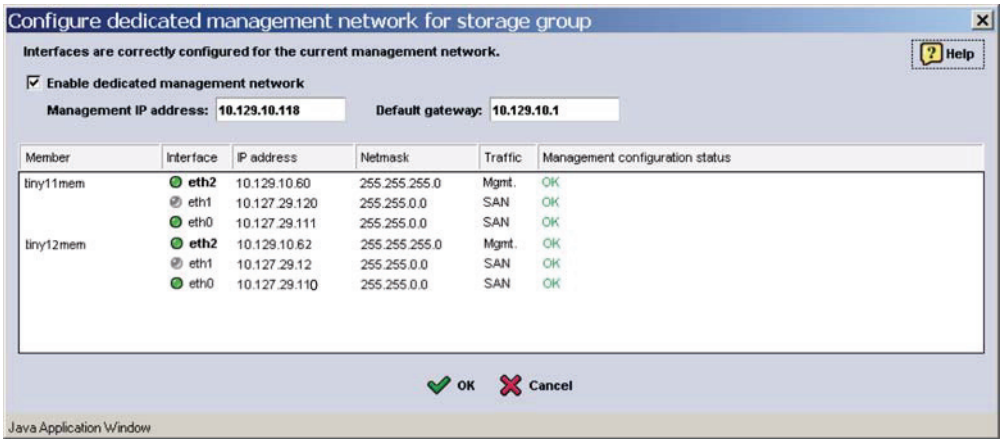

8. In the Warning dialog box, click `Yes` to restart the Group Manager GUI session using the new management IP address.

Note: 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 (Figure 4-1). The remaining interfaces for iSCSI traffic are on a different network and `SAN` appears in the `Traffic` column.

Management network post-configuration tasks

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

- Inform administrators of the new management network IP address.
- Uninstall and reinstall the Group manager GUI application If you have a shortcut to the Group Manager on the computer desktop, the group address in the shortcut is not updated with the new management address. See *Uninstalling the GUI application* on page 3-2.
- Update the group IP address in the application to the dedicated management address if you are running SAN HeadQuarters. See the SAN HeadQuarters documentation.

Displaying management network information

1. Click `Group`, then `Group Configuration`, and then the `General` tab.

   The management network address is shown in the `Management IP address` field. (If there is no management network, this field is not shown.) The General Settings panel also shows the group IP address, which you use for all iSCSI traffic, including traffic between replication partners.

2. Click `Configure management network` to display more detail.

Adding a member to a group with a management network

If you add a member to a group that has a management network, you must assign the highest-numbered network interface on the new member to the management network.

After you select the pool and RAID policy for the new member, as described in *Setting the RAID policy and pool for a new member* on page 6-6, the Modify IP Settings – Management Network dialog box appears.

In the Modify IP Settings – Management Network dialog box:

1. Enter an IP address that is on the management network.
2. Enter a subnet mask (netmask).
3. Select `Enable this interface`.
4. Select `Restrict to management access`.
5. Click `OK`.
6. Verify the configuration. See *Displaying management network information* on page 4-16.

Modifying the management network configuration

1. Click `Group`, then `Group Configuration`, and then the `Advanced tab.`
2. In the Dedicated Management Network panel, click `Configure management network`.
3. Enter the information in the `Management IP address` or `Default gateway` field.
4. In the Modify IP Settings dialog box, double-click the items and make your changes.

   Note: Make sure all the network addresses used in the management network are on the same subnet, which is different from the iSCSI network.

5. Click `OK` to confirm that you want to perform the operation. The Group Manager GUI automatically restarts, using the group IP address.

Unconfiguring a management network

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

1. Click `Group`, then `Group Configuration`, and then the `Advanced` tab.
2. In the Group Advanced window's Dedicated Management Network panel, click `Configure management network`.
3. In the Configure Management Network dialog box, deselect `Enable dedicated management network`.
4. Click `OK` to confirm that you want to perform the operation. The Group Manager GUI automatically restarts, using the group IP address.
5. Log in to the group.
6. For each member:
   a. Disconnect the former management interface from the management network.
   b. If the interface is eligible for iSCSI traffic and you want to use it, connect it to the iSCSI network.
   c. Click `Members`, then the member name, and then the `Network` tab.
   d. Select the former management interface and click `Modify IP settings`.
   e. In the Modify IP settings dialog box, if you do not want to use the interface, delete the IP address and click `OK`.

      If you want to use the interface, do the following:

      - Change the IP address and subnet mask to the iSCSI network.
      - Deselect `Restrict to management access`.
      - Select `Enable this interface`.
      - Click `OK`.
7. To log in to and manage the group, connect to the group IP address.

5 Group configuration

You can change the group configuration defaults and initial settings. You can also add members to a group to expand group capacity and improve performance.

Displaying the Group summary

The Group Summary window appears when you first start the Group Manager GUI or when you click the group name at the top of the far-left panel.

Figure 5-1: Group Summary Window

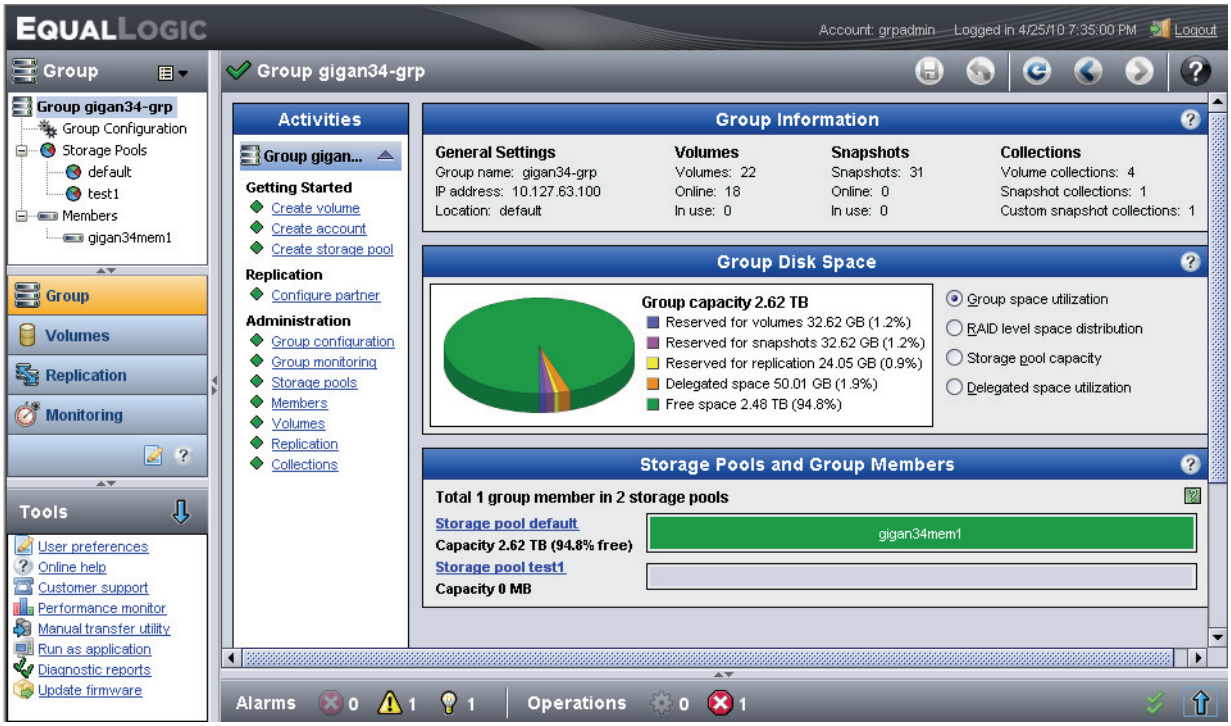

The Group Summary window displays the basic group configuration and resources in the following panels:

- Group information panel – Shows a summary of group parameters and resource use
- Group disk space panel – Shows how the group is using space. You can display details about the space by selecting from the radio button options.

  Note: The group free space might not be precise because the GUI uses a rounding algorithm to calculate free space. To display the actual free space in megabytes, you must use the CLI. Enter the `cli-settings displayinMB on` command. Then enter the `show` command to display the free group space. See the *CLI Reference* manual.

- Storage pools and Group members – Shows storage pools in the group, with details of pool capacity and used space. See *About storage pools* on page 7-1.

See the online help for information about the data fields and options.

Displaying the Group configuration

To display the current configuration, click Group Configuration in the left panel. The Summary panel contains important configuration attributes and provides links to views that contain more details and in some cases, provide options to modify group configuration attributes.

Group configuration summary panel

Table 5-1 shows the data fields available in the group configuration Summary panel and provides links to additional panels containing configuration values.

Table 5-1: Group Configuration Summary Panel

Field	Description	User Actions
General Settings	Group name and IP Address. Click the General Settings link to go to the General tab.	See *Modifying the group IP address or group name* on page 5-5
Administration Access	Whether Web access, Telnet access, or SSH access to the group are enabled or disabled. Click the Administration Access link to go to the Administration tab.	See *Enabling or disabling GUI or CLI access* on page 4-2
E-mail Notifications	Whether e-mail alerts and/or email-home functionality are enabled. Click the E-mail Notifications link to go to the Notifications tab.	*Configuring E-Mail notification* on page 14-3 *Configuring E-Mail home* on page 14-4
Event Logs	Whether syslog reporting is enabled for the group. Click the Event Logs link to go to the Notifications tab.	See *Configuring syslog notification* on page 14-5
iSCSI Authentication	Whether RADIUS and/or local CHAP authentication is enabled. Click the iSCSI Authentication link to go to the iSCSI tab.	*Configuring target authentication* on page 8-4 *Using RADIUS authentication and accounting servers* on page 4-9
SNMP Settings	Whether SNMP access and/or SNMP traps are enabled. Click the SNMP Settings link to go to the SNMP tab.	See *Configuring SNMP trap destinations* on page 14-7
VDS/VSS	Whether VDS/VSS access to the group is restricted or unrestricted. Click the VDS/VSS link to go to the VDS/VSS tab.	See *Displaying and configuring Windows service access to a Group* on page 4-12

Group configuration tabs

Each of the Group Configuration tabs provides additional information and configuration options as shown in Table 5-2:

Table 5-2: Group Configuration Tabs

Tab Name	Description	User Actions
General	General group settings, date and time.	*Modifying the group IP address or group name* on page 5-5 *Modifying the time zone and clock time* on page 5-3
Administration	Administration access, RADIUS settings, and administrative accounts.	*Administration access options* on page 4-2 *Displaying RADIUS authentication and accounting servers* on page 4-9 *Displaying local administration accounts* on page 4-5
Notifications	E-mail event notifications and the event logs.	*Configuring E-Mail notification* on page 14-3 *Configuring syslog notification* on page 14-5
iSCSI	Authentication, discovery and local CHAP accounts.	*Using CHAP accounts on a RADIUS authentication server* on page 8-3 *Configuring target authentication* on page 8-4 *Configuring the group to use an iSNS server* on page 8-5 See *Displaying local CHAP accounts* on page 8-2
SNMP	Access and traps.	*Displaying and configuring SNMP access to a Group* on page 4-11 *Configuring SNMP trap destinations* on page 14-7
VDS/VSS	Access control list.	*Displaying and configuring Windows service access to a Group* on page 4-12
Defaults	Volume settings, thin provisioning, and iSCSI settings.	*Displaying group-wide default volume settings* on page 9-5 *Displaying the iSCSI target name and alias* on page 9-2
Advanced	Load balancing, member firmware and dedicated management network.	*Enabling or disabling performance load balancing (advanced)* on page 5-6 *Updating member firmware* on page 6-13 *Configuring a management network* on page 4-14

Modifying the time zone and clock time

To display or change the current date and time values:

Click `Group Configuration`, then the `General tab`. See the online help for information about the data fields and options.

To modify the group time zone and time:

1. Select the time zone and city from the Time zone pull-down menu.
2. Click `Modify group time` in the Date and Time panel.
3. Change the date and time. To synchronize the time with the computer running the GUI, click `Set to time of system running GUI`.
4. Click `OK` to save the changes. Click `Cancel` or `Reset to current group time` to cancel the changes.

Setting the time through an NTP server

You can specify that the group use up to three external Network Time Protocol (NTP) servers to automatically set the same time for all the group members. The group uses one NTP server at a time and the first listed server is the default server. The group uses the other servers, in the order specified, if the default server is not available.

Requirement: The NTP server must be accessible to all the group members.

Recommendation: If you are using a dedicated management network, Dell recommends that the NTP server be on the same subnet as the dedicated management network.

1. Click `Group Configuration`, then `General tab`.
2. Click `Add` under NTP servers in the Date and Time panel.
3. Enter the IP address for an NTP server. Use the *`ip_address`*:*`port`* format to specify a port number other than the default, 123, then click `OK`.
4. [Optional] Repeat steps 2 and 3 to add more NTP servers. You can specify up to three NTP servers. Use the arrows to move a server up or down in the list.

Changing or deleting an NTP server

On the Date and Time panel under NTP servers, select the IP address and click `Modify` or `Delete`.

If the group time changes (which means you manually modified it or the NTP server time changed), the synchronization operation can take as long as 24 hours if the time changes by an hour or more.

Expanding group capacity

To expand PS Series group capacity and increase performance, you can add new members (arrays) to the group, with no disruption to SAN storage users:

1. Install the new array hardware and connect the array to the network. Make sure you connect the array to the iSCSI network.
2. Use the `setup` utility or the Remote Setup Wizard to configure the array as a group member. You need:
   - Group name and group IP address for the group that you want to join.
   - Group membership password.
   - IP address for a network interface on the new array (usually, Ethernet 0), in addition to the netmask and default gateway for the iSCSI network. You must connect the network interface to the network.

     See the *Installation and Setup* manual for your array model for details.
3. Perform the post-setup member tasks. See *Common group tasks*.

   Note: If you configured a dedicated management network in the group, you must configure the highest-numbered network interface on the management network. See *About dedicated management networks (advanced)*.

About group network configuration

The group network configuration, which appears in the Group Summary window, includes the group name and group IP address, which you set when creating a group:

- Group name – Identifies the group on the network. For example, when you configure a replication partner, you use the name to identify the group.
- Group IP address – 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.

Impact of modifying the group network configuration

Before modifying the group name or group IP address, make sure you understand the impact of the modification. You might need to make adjustments to your SAN environment because of the modification.

Consider the following:

- 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. See *Replication partner requirements* for information.
- 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.

Modifying the group IP address or group name

Before changing the group name or group IP address, see *Impact of modifying the group network configuration* on page 5-5.

Requirement: Group names must be unique in your network environment. A group name can consist of one to 54 letters, numbers, or hyphens. The first character must be a letter or number.

IPv6 Note: If you are using IPv6 addresses exclusively in the group, you must use the CLI to change the group IP address. See the *CLI Reference* manual for more information.

1. Click `Group`, then `Group Configuration`.
2. Select the `General` tab. (See the online help for information about the data fields and options.)
3. Modify the group name or group IP address.
4. [Optional] Change the location and group description.
5. Click `Save all changes` (Control+S).

Modifying the group membership password

To add a member to a group, you need the group membership password that is initially established when you create the group.

1. Click `Group`, then `Members`, and then `Set password.`
2. Enter the new password twice and click `OK`. The password can contain 3 to 16 alphanumeric characters and is case-sensitive.

Shutting down a group

Before shutting down a group, see *Impact of modifying the group network configuration* on page 5-5.

To perform maintenance, you might need to shut down a PS Series group. While the group is shut down, group volumes are not available.

1. Stop applications using the group volumes.
2. Disconnect iSCSI initiators from the group volumes.
3. Shut down each group member. See *Shutting down a member* on page 6-14.

   Note: Do not turn off power to a group member until you cleanly shut down the member.

To start the group, power on all group members.

Enabling or disabling performance load balancing (advanced)

Note: Dell recommends that you enable performance load balancing. Disabling load balancing degrades SAN performance.

By default, the group tries to store volume data on pool members with a RAID configuration optimal for volume performance, based on metrics collected from each group member. You can enable or disable automatic performance load balancing. Other load balancing methods continue to apply.

1. Click `Group`, then `Group Configuration`, and then the Advanced tab.
2. In the Load Balancing panel, select or deselect `Enable performance load balancing in pools` (Alt+E).
3. Click `Save all changes` (Control+S).

6 Group members

A PS Series group includes one or more PS Series arrays configured as group members.

Displaying Group members

To display information about all group members, click `Members` in the far-left panel. The Group Members panel appears, containing the following panels:

- Group disk space panel – Provides information about capacity use for the group and pools and capacity use distribution by raid level.
- Group members panel – Provides information about individual members (arrays) in the group, physical attributes, capacity and capacity use.

See the online help for information about the data fields and options.

Displaying member details

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the appropriate information tab.

You can display information about a group member as follows:

- Member status tab – The general status of a member, such as member information and hardware health.
- Enclosure tab – Array hardware status.
- Controllers tab – The status of controllers, such as which controller is active.
- Disks tab – The status of disks and model or revision of firmware.
- Network tab – Ethernet interface operational status.
- Connections tab – iSCSI connections to storage objects such as volumes.
- Service tab – Information, such as details of the type, model and firmware revision of components in the array.

See the online help for information about the data fields and options.

Member status tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Status` tab.

The following panels provide information about the member:

- General member information – Provides information about the array model, its configuration settings, and RAID status.
- The Member Health Status panel – Provides a physical view of the array hardware and enables you examine the following details:

  - Click `Front view` to see the front panel of the array. For some array models, the GUI shows disk drives located behind the front bezel.
  - Click `Inside view` (not available on all array models) to see the disk drives located inside the array.
  - Click `Rear view` to see the back panel of the array, including the control modules and the power supply and cooling modules. The front and rear views shown in your GUI depend on the array model of the group member.

    - A red “X” over a hardware component indicates uninstalled or unconfigured hardware. Place the pointer over a component to display status details.

    - The Member Health Status panel also shows a table with the alarm status for the array hardware components.
  - Click `View Alarms` to display all the alarms in the member.

- Member space panel – Provides information about the member’s storage capacity and how storage space is allocated. Be aware of the following:
  - Total member capacity does not include the space in spare disk drives because the member uses spare disk drives only if a failure occurs). Member capacity depends on the number and size of the installed disks and the selected RAID policy.
  - If free member space is low, you can increase the member capacity by adding disk drives, if empty slots are available, or by replacing the current drives with higher-capacity drives. You can also add more group members to increase overall group capacity.

See the online help for information about the data fields and options.

Enclosure tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Members` tab.

The following panels provide information about the member’s enclosure (array enclosure):

- Power supplies panel – Enables you to identify power supply locations and verify their status.
- Cooling fans panel – Enables you to identify cooling fan locations and verify their status, including operational parameters.
- Temperature sensors panel – Enables you to identify temperature sensor locations, verify their status and verify operational parameters.

See the online help for information about the data fields and options. Your array hardware manual contains more information about physical features of the array.

See also *Monitoring the member enclosure* on page 15-17 and *Monitoring cooling and fans* on page 15-18.

Controllers tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Controllers` tab.

The following panels provide information about the member's controller modules:

- Control module panel – Enables you to identify controller locations and verify their status and cache battery condition. Identifies the current PS Series firmware
- Memory cache panel – Describes the cache mode and current caching policies.

See the online help for information about the data fields and options. See also *Monitoring control modules* on page 15-20.

Disks tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Disks` tab.

The following panels provide information about the member's disks:

- Disk array summary panel – Enables you to identify disk drive locations and verify the status of a disk. Provides information about failed disks and operations in progress.

  Note: The disk slots are color-coded to represent the status. Mouse over a disk to view its status information as text.

- Installed disks panel – Describes the physical attributes of installed disks. You can select an individual disk and display disk activity statistics and performance data. See *Using the Performance Monitor* on page 15-29.

See the online help for information about the data fields and options. See also *Monitoring disk drives* on page 15-22.

Network tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Network` tab.

The following panels provide information about the member's network interfaces:

- Status of network interfaces panel – Provides a description of the operational status of network interfaces in the member's controllers, and the physical attributes of the interface. Enables
- IP configuration panel – Describes the IP configuration values for each network interface and enables you to modify the IP configuration.

See the online help for information about the data fields and options. See also *About member network configuration* on page 6-7.

Connections tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Connections` tab.

The iSCSI Connection panel displays details of the current initiator connections. See the online help for information about the data fields and options. See also *Monitoring iSCSI connections* on page 15-3.

Service tab

1. Click `Group`, then expand `Members`.
2. Select the member name, and click the `Service` tab.

The following panels provide information about the member's hardware components:

- Component versions panel – Provides information that enables you to identify your array and its components for the purpose of service and upgrades, including installed firmware revisions.
- Disk versions panel – Provides information that enables you to identify disk drives for the purpose of service and upgrades, including installed firmware revisions.

You can also restart or shut down a member from this panel, as described in *Restarting a member* on page 6-15 and *Shutting down a member* on page 6-14

See the online help for information about the data fields and options. See also *Displaying member service information*.

Member RAID policies

PS Series arrays protect data by using RAID technology and spare drives. After you add a member to a PS Series group, you choose the RAID policy for that member.

To display current RAID policies, click `Members` in the far-left panel. The Group Members panel shows the RAID policy for each member.

Supported RAID policies are:

- RAID 10
- RAID 10 no spares
- RAID 50
- RAID 50 no spares
- RAID 5
- RAID 6
- RAID 6 no spares
- RAID 6 Accelerated (supported only on array models with disk drive configurations that include solid-state drives and hard disk drives.)

The RAID policy for a member consists of two parts:

- RAID level – RAID 10, RAID 50, RAID 5, or RAID 6. See *RAID level characteristics*.

  Recommendation: For optimal performance, Dell recommends that you assign the same RAID level to pool members with the same disk type and disk speed.

- Spare drive policy – Whether the member automatically configures and uses spare disk drives. Spare drives increase availability.

  The number of spare drives depends on the array model and the number of installed drives. For RAID 6 Accelerated, only one hard disk drive is configured as a spare. The solid-state drives are not protected by sparing.

  Recommendation: Dell recommends that you use a spare-drive RAID policy. Only use a no-spare-drive RAID policy if you have sufficient support staff and maintain a stock of replacement disk drives.

The storage in the member is available after you set the RAID policy. The member automatically configures the disk drives according to the designated RAID level, with the appropriate number of spare drives.

Although all RAID levels provide good performance and data protection, there are some differences. When choosing a RAID policy, you should identify the performance and availability needs of your workload and select a RAID policy that meets those needs. If your workload has mixed requirements in terms of performance and availability, you might want to consider mixing RAID levels in a multi-member group.

RAID level characteristics

The characteristics of each supported RAID level are:

- RAID 10 – Striping on top of multiple RAID 1 (mirrored) sets.
- RAID 50 – Striping on top of multiple RAID 5 (distributed parity) sets.
- RAID 5 – One or more RAID 5 sets.
- RAID 6 – One or more RAID 6 (dual parity) sets.

  RAID 6 Accelerated is supported only on array models with disk drive configurations that include both solid-state drives and hard disk drives. RAID 6 Accelerated optimizes the use of solid-state drives for critical data. One hard disk drive is configured as a spare and provides redundancy protection in the event of a hard disk drive failure or a solid-state disk drive failure.

Each RAID 1 set or RAID 5 set can survive a single disk drive failure. A RAID 6 set can survive two simultaneous drive failures. If a drive fails in a RAID set, the RAID set is degraded.

Consider the following performance and availability factors when choosing a RAID level for a member:

- RAID 10 and RAID 50 provide excellent reliability for your data, in addition to overall high performance.
- RAID 10 provides the best performance for workloads that are mostly small random writes.
- RAID 6 provides high availability, but at the expense of performance during data reconstruction.
- RAID 6 is not recommended for workloads consisting mainly of random writes.

Table 6-1 compares the performance and availability characteristics of the supported RAID levels. The first column lists workload requirements, with the other columns respectively listing the performance quality for each requirement at RAID 10, RAID 50, RAID 5, and RAID 6.

Table 6-1: RAID Level Characteristic Comparison

Workload Requirement	RAID10	RAID 50	RAID 5	RAID 6
Capacity	Average	Good	Excellent	Good
Availability	Excellent	Good	Average	Excellent
Sequential reads	Excellent	Excellent	Excellent	Excellent
Sequential writes	Good	Good	Good	Good
Random reads	Excellent	Excellent	Excellent	Excellent
Random writes	Excellent	Good	Average	Average
Performance impact of drive failure or RAID reconstruction	Minimal	Moderate	Moderate to heavy	Heavy

Supported RAID policy conversions

While a member remains online, you can convert it from one RAID policy to another *only* if the new RAID policy provides the same or more space than the current policy.

Table 6-2 shows the supported RAID policy conversions. The first column lists RAID configurations, and the second column lists supported RAID policy conversions for them.

Table 6-2: Supported RAID Policy Conversions

Current RAID Policy	Supported Conversion
RAID 10	All
RAID 10 no spares	RAID-5, RAID-50, RAID-50 no spares, RAID-6, RAID6 no spares
RAID 50	RAID-5, RAID-50 no spares, RAID-6, RAID6 no spares
RAID 50 no spares	RAID-5,RAID-6, RAID6 no spares
RAID 5	None
RAID 6	RAID-5
RAID 6 no spares	RAID-5
RAID 6 Accelerated	None

If a RAID policy conversion is not supported, you can remove the member from the group and then add it to the group again. You can then set the RAID policy.

Setting 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. See *Member RAID policies*.

Notes: 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.

If you want to use a no-spare disks RAID policy, you must use the Group Manager CLI to set the RAID policy.

1. In the Group Summary window, expand `Members` and double-click the member name or click `Group`, then expand `Members`, and then select the member name.

   The GUI shows whether a member is configured or not.

2. In the warning dialog box that appears, click `Yes` to configure RAID on the member.

3. In the Configure Member – General Settings dialog box, select the pool and click `Next`.

4. If prompted, confirm you want to assign the member to the pool.

   Recommendation: For optimal performance, Dell recommends that you assign the same RAID level to pool members with the same disk type and disk speed.

5. In the Configure Member – RAID Configuration dialog box:

   - Select the RAID policy.
   - By default, member storage space is immediately available, although performance is not optimal until the RAID verification completes. To make space unavailable until the RAID verification completes and batteries are fully charged, select `Wait until the member storage initialization completes`.

6. Click `Next`.

7. When the configuration is complete, click `Finish` in the Configure Member – Summary dialog box.

Converting a RAID policy

See *Supported RAID policy conversions*.

Note: To convert to a no-spare-drive RAID policy, use the Group Manager CLI.

1. Click `Group`, then expand `Members`, then select the member name, and then click `Modify RAID configuration`.

2. In the Modify RAID Configuration dialog box, select the new RAID policy.

   The values in the Member Capacity table automatically update, based on the RAID policy you selected.

3. To ensure member space is not available until the RAID verification completes, you can select `Wait until the member storage initialization completes`. If this option is already selected, you cannot change the selection.

4. Click `OK`.

   While the RAID policy is changing, the member's RAID status is `expanding`.

About member network configuration

To enable network-based management, group communication, and iSCSI traffic, each group member must have at least one functioning network interface that you connected to the network and configured with an IP address and subnet mask (netmask).

Recommendation: For high availability and performance, Dell recommends that you configure multiple network interfaces and use redundant network switches.

Note: Some control module types include a network interface that can be used only in a dedicated management network.

Member network requirements and recommendations

The minimum network requirement for a group member is one functioning network interface with the following characteristics:

- Located on the active control module
- Connected to the network
- Configured with an IP address and netmask (at least one network interface on a member must be on the same subnet as the group IP address)
- Enabled

When you add a member to a group, you configure a network interface with an IP address and netmask and enable the interface (typically, Ethernet 0).

Warning: The minimum network configuration is a single-point-of-failure configuration. A network interface failure, control module failure, or network or switch failure causes the member and any volumes with data to go offline until you correct the problem. If it is the only member in the group, the group becomes inaccessible from the network.

To increase performance and availability, Dell recommends that you expand the minimum network configuration as follows:

- Configure redundant network connections. See *Configuring redundant network connections*.
- Configure redundant control modules. See *Configuring redundant control modules*.
- Configure redundant network switches. See *Configuring redundant network switches*.

See the *Hardware Maintenance* manual for your array model for network configuration details.

Configuring redundant network connections

Redundant network connections protect against network interface failure and increase performance (network bandwidth). If a network interface fails or you disconnect it, another interface can continue to service I/O requests.

1. Connect two or more network interfaces on the active control module to the network.
2. Configure the interfaces (assign an IP address and netmask and enable each interface). See *Configuring network interfaces*.

Configuring redundant control modules

Redundant control modules protect against control module failure by enabling control module failover. If the active control module fails, the secondary control module takes over and becomes active.

1. Install a secondary control module in the member.
2. For each configured network interface, connect the Ethernet port on the secondary control module to the network.

Configuring redundant network switches

Redundant network switches protect against network and switch failures and improve network performance.

Requirement: Requires redundant network connections. See *Configuring redundant network connections*.

1. Distribute redundant network connections across multiple switches.
2. Connect switches to links with sufficient inter-switch bandwidth.

Displaying the member network configuration

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Network` tab.
2. In the Member Network window, review the current network configuration for the member.

Configuring network interfaces

When you add a member to a group, you configure only one network interface. It is best practice to configure all network interfaces eligible for iSCSI traffic. The number of iSCSI eligible interfaces depends on the type of control module.

You can also change the existing member network configuration or enable or disable an interface. If you change the IP address for a functioning network interface, the member disconnects all iSCSI initiators from that interface. Most initiators reconnect automatically.

IPv6 Note: If you are using IPv6 addresses exclusively in the group, you must use the CLI to configure a network interface. See the *CLI Reference* manual.

1. Connect the network interface on the active control module to the network.

   Recommendation: Dell recommends that you also connect the network interface on the secondary control module to increase availability.

2. Obtain the IP address and netmask for the interface. At least one network interface on a member must be on the same subnet as the group IP address.
3. Click `Group`, then expand `Members`, then select the member name, and then click the `Network` tab.
4. In the IP Configuration panel, select a network interface and click `Modify IP settings`.
5. Enter the following information in the Modify IP Settings dialog box:
   - IP address for the network interface.
   - Subnet mask (netmask) for the network interface IP address.

Note: Unless you are using a dedicated management network, the default gateway is the same for all network interfaces on an array. To modify the default gateway, see *Modifying the default gateway for a member*.

6. Select `Enable this interface`. You must enable the interface in order to use it.
7. Click `OK`.

Enabling or disabling a network interface

Enabling a network interface makes it operational if you configure it properly. Disabling a network interface makes it unable to service network I/O requests, but does not unconfigure the interface.

Disabling a functioning network interface disconnects all iSCSI initiators associated with the member. Most initiators reconnect to another interface automatically.

Requirement: A member must have at least one functioning network interface that is configured, enabled, and connected to a network. If you disable the only functioning network interface on a member, the group sets the member offline. If the member is the only member in the group, the group is not accessible from the network.

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Network` tab.
2. In the IP Configuration panel, select the network interface.
3. In the Activities panel, click `Enable interface` to enable the interface or click `Disable interface` to disable the interface.

Unconfiguring a network interface

If you unconfigure a functioning network interface, the member disconnects all iSCSI initiators from that interface. Most initiators reconnect to another interface automatically.

Requirement: A member must have at least one network interface that is functioning, configured, enabled, and connected to a network. If you unconfigure the only functioning network interface on a member, the group sets the member offline. If the member is the only member in the group, the group becomes inaccessible from the network.

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Network` tab.
2. In the IP Configuration panel, select the network interface.
3. In the Activities panel, click `Modify IP Settings.`
4. In the `IP address` field, delete the IP address.
5. Click `OK`.

Modifying the default gateway for a member

When you add a member to a group, you can specify a default gateway that the member uses for all its network interfaces. You can later modify this setting.

Requirement: You must use a default gateway to enable communication outside the local network.

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Network` tab.
2. In the IP Configuration panel next to the `Default gateway` field, click `Modify`.
3. Specify a default gateway IP address if you want to use a default gateway, or delete the IP address if you do not want to use a default gateway.
4. Click `OK`.

Modifying a member name or description

Requirement: Member names must be unique in a group. A member name can consist of one to 63 letters, numbers, or hyphens. The first character must be a letter or number. A member description can be up to 127 alphanumeric characters.

1. Click `Group`, then expand `Members`, then select the member name, and then click `Modify member settings.`
2. In the Modify Member Settings dialog box, modify the member name or specify a description for the member.
3. Click `OK`.

About write cache operations

Each (active) control module contains a battery-backed cache. The active control module cache operates in one of two modes:

- Write-back – The member stores data in the cache until it is written to disk drives, which can improve performance. The cache battery protects the data in the cache. In a dual control module configuration, the member mirrors cache data across the two controllers, providing additional data protection.
- Write-through – The member immediately writes data to the disk drives, which might decrease performance.

You can set policies that control the mode when the battery is low or a control module fails in a dual control configuration. See *Setting write cache policies*.

Setting write cache policies

You can set policies that control the mode when the battery is low or a control module fails in a dual control configuration.

Recommendation: Dell recommends that you keep the default write cache policy settings for the best performance and availability.

- Use write-through mode if only one controller is functional – Also called single-controller-safe mode, if you enable this policy on a member with a single control module, the cache always uses write-through mode. If you disable this policy (the default), under non-failure conditions, the cache uses write-back mode.

  If you enable this policy on a member with dual control modules, but only one module is functioning, the cache

on the functional control module uses write-through mode. If you disable this policy (the default), the cache on the functioning control module uses write-back mode.

- Use write-through mode if battery charge is below tolerance – Also called low-battery-safe mode, if you enable this policy (the default) on a member with a single control module, the cache uses write-through mode if the charge on the cache battery is low.

  If you enable this policy on a member with dual control modules, but do not enable the single-controller-safe policy, the active control module cache uses write-back mode if its cache battery has an adequate charge, but the secondary control module's cache battery has a low battery charge. If you enable the low-battery-safe policy and the single-controller-safe policy, the active control module cache uses write-through mode if the battery charge on either control module is low.

  If you disable the low-battery-safe policy (not recommended), a cache uses write-back mode, regardless of cache battery charge.

To change the write cache policies, see *Modifying write cache policies*.

Modifying write cache policies

Recommendation: Dell recommends that you keep the default write cache policy settings for the best performance and availability:

- Do not select (disable) `Use write-through mode if only one controller is functional`.
- Select (enable) `Use write-through mode if battery charge is below tolerance`.

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Controllers` tab to display the Member Controllers window.
2. In the Memory Cache panel, select or deselect the cache mode policies.
3. Click `Save all changes` (Control+S).

About member firmware

Each control module you install in a group member must be running the same version of the PS Series firmware. Firmware is stored on a compact flash card or a MicroSD card on each control module.

Recommendation: Dell recommends that you always run the latest firmware to take advantage of new features and fixes.

Dell recommends that all group members run the same firmware version. If you are adding a new array to a group, update the group to the latest firmware before adding the new member.

You can upgrade member firmware to a higher version or downgrade member firmware to a lower version.

There are two methods for updating firmware:

- `update` command – See the *CLI Reference* manual.
- Group Manager GUI – See *Updating member firmware*.

Before updating firmware, see *Firmware update considerations and prerequisites* for important information to consider.

Firmware update considerations and prerequisites

Regular firmware updates are an important part of maintaining a well-functioning group.

Before performing any firmware update, read the *Release Notes* for the new firmware and the *Updating Storage Array Firmware* document. In addition, see your PS Series support provider for detailed information about firmware and firmware updates.

When updating array firmware, keep in mind the issues listed below. The *Updating Storage Array Firmware* document describes them in greater detail and provides recommendations for addressing them.

- **SAN planning.** Before updating a SAN component, you must fully understand the impact of the update on the infrastructure. Careful planning of the upgrade process can help you avoid unplanned downtime.
- **Timing.** Dell recommends that you perform firmware upgrades during off hours or scheduled maintenance periods to avoid disruption of service to the applications and servers that the storage group supports. When scheduling a firmware update, allow enough time to update and restart the entire group of arrays.
- **Planning for downtime and minimizing host disruption.** The update procedure requires an array restart. During an array restart, volumes with data on the array are temporarily unavailable until the restart completes. To make sure that applications are not affected, follow the Dell guidelines in the *iSCSI Initiator and Operating System Considerations* document.
- **Backing up data.** Upgrades should be implemented after a backup has occurred.
- **Working with Multi-Member Groups.** Dell recommends that all PS Series group members run the same version of the storage array firmware. The PS Series Release Notes describe which firmware versions can co-exist in a group; however, only those features and bug fixes common to all versions are available.
- **Software Prerequisites.** Depending on the firmware version, you might be required to update certain software components prior to applying the upgrade.
- **Array Prerequisites.** Before upgrading the firmware, you need network access to all group members, the group IP address, and access to the `grpadmin` account.
- **Supported Upgrade Paths.** Usually, you can update an array directly to the latest firmware version. However, in some cases, you might need to update to an interim version before updating to the latest version.

Updating member firmware

You can upgrade PS Series firmware to a higher version or downgrade firmware to a lower version.

Procedure for updating firmware

The *Updating Storage Array Firmware* document provides step-by-step instructions for updating member firmware. You can obtain this document from the downloads area of the EqualLogic Customer Support Website.

Disallowing member firmware downgrades

If the Disallow downgrades option is active, the group is not using all features of the installed firmware. You must disallow firmware downgrades to use the features in an updated firmware release.

Requirement: Before you disallow firmware downgrades, make sure all group members are running the same firmware version.

1. Click `Group`, then `Group Configuration`, and then the `Advanced` tab.
2. Click `Disallow downgrades` in the Member Firmware panel. (Alt+D).
3. Confirm that you want to disallow downgrades.

Removing a member from the group

You can remove a member from a multi-member group while the group is online and data stays available. The group moves any volume data from the member you are removing to the remaining pool members. Removing a member decreases the overall storage capacity of the group and also decreases the capacity of the storage pool to which the member belongs.

Note: You cannot remove a member from the group unless the remaining pool members have enough space to store the data from the member that you want to remove.

Members you remove from a group are automatically reset to the factory defaults. Previous group, member, and volume information, as well as any volume data on the member, are removed.

Removing a member from a group can be a long operation, depending on the amount of data the group must move to the remaining pool members. While data is moving, the member status is `vacating-in-progress`.

1. Click `Group`, then expand `Members`, then select the member name, and then click `Delete member.`

   If you receive a “vacate failed” message, the member is offline.
2. Confirm you want to delete the member.

Shutting down a member

For maintenance purposes, you might need to cleanly shut down a member. Shutting down a member has no effect on member, volume, or group configuration information or volume data stored on the member.

Shutting down a member does not turn off array power. To turn off array power, turn off all power switches on the array after the shutdown completes.

Note: To restart a member that you shut down, turn on all power switches.

1. Click `Group`, then expand `Members,` then select the member name, and then click the `Service` tab.
2. Click `Shutdown` in the lower part of the GUI window.
3. In the Member Shutdown confirmation dialog box, enter your group administrator account password and click `OK`.

The group sets any volumes with data on the member offline. The group sets the volumes online when you restart the member.

Restarting a member

When you restart a member, the group sets any volumes with data on the member offline. The group sets the volumes online when the restart completes. Restarting a member has no effect on member, volume, or group configuration information or volume data stored on the member.

Restriction: Do not repeatedly restart a member.

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Service` tab.
2. In the Member Services window, click `Restart.`
3. In the Member Restart confirmation dialog box, enter your administration account password.
4. Click `OK`.

Part II: Using Group Storage Space

7 Storage pools

About storage pools

Storage pools allocate storage space into partitions comprising one or more members.

By default, a group provides a single pool of storage. If your group has multiple members, you can divide group space into different pools and then assign members.

Note: Load balancing operates only within pools.

By default, a PS Series group provides a single pool of storage. From this pool, you allocate space to users and applications by creating volumes, which are seen on the network as iSCSI targets.

When you first create a group, there is one storage pool in a group, called `default`. Unless you specify a different pool, members and volumes are assigned to the default pool.

Restriction: You cannot delete the default storage pool.

However, some environments might need to segregate storage space. In this case, you can divide PS Series group space into multiple storage pools. Using this "SAN within a SAN" technology, administrators can easily separate workloads, while retaining the advantages of storage consolidation.

For example, you can segregate storage space according to application, service level, RAID type, or department. Mission critical applications can use storage resources that ensure fast and consistent performance, while little-used or archived data can use different resources.

To use multiple storage pools, a group must contain more than one member. You can then assign members to different pools. A pool can contain multiple members. Performance load balancing occurs only across the pool members. You can create up to four pools in a group.

After you assign a member to a pool, you can assign volumes to the pool. Only the resources that the pool members provide are available to the volumes. As capacity needs change, you can move members or volumes from one pool to another, while data remains online.

For detailed information about using storage pools, see the Technical Report *Deploying Pools and Tiered Storage in a PS Series SAN*, which you can download from the EqualLogic customer support website.

Configuring a storage pool

To plan a storage pool configuration, you must understand the capacity and RAID policy of your group members and also the capacity, performance, and service level needs of your volumes. This enables you to select the appropriate members for each pool and then identify the volumes to assign to the pools. Try to identify future demands for capacity and performance when planning the storage pool configuration.

Identify the following:

- Capacity and performance of each member. Consider the following factors:
  - Disk type, either Serial Attached SCSI (SAS) or Serial Advanced Technology Attachment (SATA)
  - Disk size

  - Disk speed
  - Member RAID level (RAID 10, RAID 50, RAID 5, RAID 6)

Recommendation: Dell recommends that pool members with the same disk spin rate have the same RAID level. For example, if a pool contains two members that have 7200 RPM disks installed, configure both members with the same RAID level.

- Disk space and performance needs of each application. Identify which applications deserve priority and calculate their disk space and network bandwidth needs and performance characteristics. For example, some applications do many random data transfers, while others do a small number of large sequential data transfers. In addition, identify rarely-accessed volumes that you use only to archive data.
- How the pools are organized. Identify the number of pools and the members that belong to each pool, based on the member performance characteristics and capacities and the requirements of the volumes. For example, you can organize your pools according to disk speed. You could also organize pools according to member RAID level.
- Which volumes belong to each pool. Put volumes with the highest service level requirements in a pool whose members can provide that service level. If an application uses multiple volumes, those volumes can be in different pools.

  When deciding which volumes belong to a pool, consider the RAID level of the pool members. See *RAID level characteristics* on page 6-5. If you do not know the optimal RAID level for a volume, make sure that the pool to which the volume belongs contains members with a variety of RAID levels, so automatic performance load balancing can occur.

Creating a storage pool

Prerequisite: See *Configuring a storage pool* on page 7-1 to understand how to select the appropriate members for each pool and then identify the volumes to assign to the pools.

There are two ways to create a storage pool:

- Create an empty pool. You can then assign members to the pool.
- Create a new pool and immediately move a member from its current pool to the new pool.

  When you move a member from one pool to another, the remaining pool members in the original pool must enough space to store any volume data that is on the member. Moving a member to a different pool can take a long time, depending on the amount of data that the group must move.

Requirement: A storage pool name can be up to 63 characters. Valid characters include letters, numbers, and hyphens. The first character must be a letter or number. A pool description can be up to 127 alphanumeric characters.

To create an empty pool:

1. Click `Group`, then `Storage Pools`, and then `Create storage pool.`
2. In the Create Storage Pool dialog box, specify a pool name and description and click `OK`.

To create a new pool and immediately move a member to the pool:

1. Click `Group`, then expand `Members`, then select the member name, and then click `Modify member settings.`

2. Under Storage pool assignment, click `Create new pool`.
3. In the Create Storage Pool dialog box, enter a name and description for the new pool and click `OK`. If a description contains spaces, surround it with quotation marks.
4. In the Modify Member Settings dialog box, click `OK`.
5. Confirm that you want to create the pool.

   The member status shows as `moving` until the move operation completes.

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, click `Members`, then *`member_name`*, then `Cancel member move.`

Displaying storage pools

To display the storage pools in a group, click `Group` and then `Storage Pools`. The Storage Pool Summary window appears, containing the following panels:

- Group disk space panel – Shows available storage space and space usage. Provides information about RAID distribution. You can change the view by selecting an option.

  Recommendation: Dell recommends that free pool space does not fall below the following, whichever is smaller:

  – 5% of the total pool space

  – 100 GB multiplied by the number of pool members

- Storage pools panel – Shows pool configuration details. Provides information about storage objects in the pool, such as volumes.

See the online help for information about the data fields and options.

Displaying storage pool details

To display the details of an individual storage pool, click `Group`, then expand `Storage Pools`, and then select the pool name.

The Status, Volumes, and Volume Templates tabs provide storage pool information and links to task activities.

Storage pool status tab

The Status tab provides general pool information, disk space information, and details of each PS Series group member in the pool. It contains the following panels:

- General pool information panel – Provides information about members, volumes and snapshots.
- Pool disk space panel – Provides information about available storage space, space usage, and RAID distribution. Enables you to select a volume administrator account.

- Pool members panel– Provides information that enables you to identify members in the pool, member status and configuration, and pool iSCSI connections.

See the online help for information about the data fields and options.

Storage pool volumes tab

To display the volumes in a storage pool: click `Group`, then expand `Storage Pools`, then select the pool name, and then click the `Volumes` tab.

The Pool Volumes panel provides a table of information about:

- Volume name, space allocation, reserved space, and accessibility status.
- Replication partners
- Snapshots for the volume
- iSCSI connections to the volume.

The table is sorted alphabetically by Volume name. You can sort the information in this panel table by clicking the heading of any column.

Moving a member to a pool

When you add a member to a group, you assign the member to a storage pool. You can also move a member from its current pool to different pool, with no effect on users or data availability.

Moving a member from one pool to another decreases the capacity of the current pool and increase the capacity of the destination pool.

Requirement: To move a member from one pool to another, the remaining pool members must have enough free space to store the data from the moved member. Moving a member to a different pool can take a long time, depending on the amount of data that the group must move to the remaining current pool members.

1. Click `Group`, then expand `Members`, then select the member name, and then click `Modify member settings`.
2. In the Modify Member Settings dialog box, select the pool for the member.

   The group updates the Pool Space table to show the new amounts of free and used space in each pool. If a pool does not have enough space to accommodate the pool change, the table cell showing free pool space displays a negative value.
3. Click `OK`.

The member status is shown as `moving` until the move operation completes.

If necessary, you can cancel an in-progress member pool move operation. The member immediately moves to the original pool.

To cancel an in-progress member pool move operation, click `Members`, then *`member_name`*, then `Cancel member move.`

Moving a volume to a pool

When you create a volume, you assign the volume to a pool. The group stores volume data on the pool members. You can also move a volume from its current pool to a different pool, with no effect on users or applications.

While a volume is moving, the group allocates volume space in the current pool and in the new pool.

Restriction: Thin clones inherit the pool setting of the template volume. You cannot move a thin clone separately from its template column.

Requirement: The pool to which you want to move a volume must have free space equal to the volume reserve and any snapshot reserve and local replication reserve for the volume.

1. Click `Volumes`, then expand `Volumes`, then select the volume name, then click `Modify volume settings`, and then click the `General` tab.
2. In the Modify volume settings dialog box, select the destination pool for the volume. The group updates the Pool Space table to show the new amounts of free and used space in each pool. If a pool does not have enough space to accommodate the pool change, the table cell showing free pool space displays a negative value.
3. Click `OK`.

The volume status is shown as `moving` until the pool move operation completes.

If necessary, you can cancel an in-progress volume pool move operation. Volume data moves to the original storage pool.

To cancel an in-progress volume pool move operation, click `Volumes`, then *`volume_name`*, then `Cancel volume moving`.

Merging storage pools

You can merge any storage pool except the default pool into another pool, called the destination pool. Merging a pool into a destination pool moves the pool members and volumes data into the destination pool. The group then deletes the empty pool.

1. Click `Group`, then expand `Storage pools`, then select the pool name, and then click `Merge storage pool.`
2. In the Merge Storage Pools dialog box, select the destination pool and click `OK`.

Modifying a storage pool name or description

You can change the name or description of a storage pool, including the default pool.

Requirement: A storage pool name can be up to 63 characters. Valid characters include letters, numbers, and hyphens. The first character must be a letter or number. A pool description can be up to 127 alphanumeric characters.

1. Click `Group`, then expand `Storage pools`, then select the pool name, and then click `Modify pool settings`.
2. Modify the pool name or description and click `OK`.

Deleting a storage pool

If you delete a storage pool, the group immediately moves its members and volumes to the default pool.

Note: You cannot delete the default pool.

1. Click `Group`, then expand `Storage pools`, then select the pool name, and then click `Delete storage pool`.
2. Confirm that you want to delete the pool.

8 iSCSI target security

Volumes and snapshots are seen on the network as iSCSI targets. It is important to understand how to protect your volumes and snapshots from unauthorized and uncontrolled access by iSCSI initiators.

About 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 8-1.

Table 8-1: 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 configured in the initiator. These credentials apply to all group targets. See *Configuring target authentication* on page 8-4.
Target access controls	The initiator must meet all the conditions in one access control record for the target. See *About iSCSI target access controls* on page 8-1.

About iSCSI target access controls

PS Series groups use access control records to prevent unauthorized computer access to iSCSI targets (volumes or snapshots).

A volume and its snapshots share a list of access control records (up to 16 for each volume). An access control record can apply to the volume, its snapshots, or both. For example, you can let a computer have access to the volume and its snapshots or access only to the volume.

When you create a volume, you can set up one access control record. You can later set up additional records.

To log in to a volume or snapshot, the initiator must comply with conditions specified in one access control record. You can specify one or more of the following conditions:

- IP address – Restricts access to iSCSI initiators that match the specified IP address (for example, 12.16.22.123). Use asterisks to indicate that any value is accepted in an octet (for example, 12.16.*.*).
- iSCSI initiator name – Restricts access to iSCSI initiators that match the specified name (for example, `iqn.2000-05.com.qlogic.qla-4000.sn00044)`.
- CHAP user name – Restricts access to computers that supply the specified CHAP user name and its associated password (or “secret”). The credentials must match a local CHAP account or a CHAP account on an external RADIUS server. See *Authenticating initiators through CHAP*.

For example, if a volume has only one access control record, which includes an IP address and CHAP user name, only a computer with that IP address and the appropriate CHAP credentials can access the volume. If an administrator creates another record that includes an iSCSI initiator name, a computer with that initiator name can also access the volume.

You can create an access control record that gives unlimited computer access. However, Dell does *not* recommend unlimited computer access unless you are testing access to a target.

Authenticating initiators through CHAP

CHAP 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 record. 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"). There are two options 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. See *Displaying local CHAP accounts* on page 8-2.
- CHAP accounts on an external RADIUS authentication server. Using a RADIUS server to manage CHAP accounts is beneficial if you are managing a large number of accounts. However, computer access to targets depends on the availability of the RADIUS server. See *Using CHAP accounts on a RADIUS authentication server* on page 8-3.

Note: If you use CHAP for initiator authentication, you can also use target authentication for mutual authentication, which provides additional security. See *Configuring target authentication*.

Displaying local CHAP accounts

To display local CHAP accounts, click `Group`, then `Group Configuration`, and then the `iSCSI` tab.

The Local CHAP accounts panel provides information about the account credentials, account status, and the related administration account. See the online help for information about the data fields and options.

Creating a local CHAP account

To create a local CHAP account:

1. Click `Group`, then `Group Configuration`, and then the `iSCSI` tab. The Group Configuration – iSCSI window appears.
2. Optionally, in the iSCSI Authentication panel, select `Consult locally defined CHAP accounts first`. If selected, credentials that an iSCSI initiator supplies are checked against local CHAP accounts before external CHAP accounts on a RADIUS server.
3. In the Local CHAP Accounts panel, click `Add`.
4. In the Add CHAP Account dialog box:
    - Enter a CHAP user name and (optionally) a password. If you do not enter a password, the group automatically generates a password that is 16 ASCII characters in length.

Note: 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.

- Select whether to enable the account. You must enable an account to use it for initiator authentication. You can later modify an account and enable or disable it.
- Click `OK`.

5. In the Group iSCSI window, click `Save all changes` (Control+S).

After creating the CHAP account, you can create an access control record and use the CHAP user name in the record. See *Configuring access control records*.

If you want to enable target authentication (for mutual authentication), see *Configuring target authentication* on page 8-4.

Modifying a local CHAP account

1. Click `Group`, then `Group Configuration`, and then the `iSCSI` tab. The Group Configuration – iSCSI window appears.
2. Select the account name in the Local CHAP Accounts panel and click `Modify`.
3. Change the name or password or enable or disable the account.
4. Click `OK`.

Deleting a local CHAP account

1. Click `Group`, then `Group Configuration`, then the `iSCSI` tab. The Group Configuration – iSCSI window appears.
2. Select the account name in the Local CHAP Accounts panel.
3. Click `Delete`.

Using 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. See the prerequisites in *Using RADIUS authentication and accounting servers* on page 4-9.

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

2. Click `Group`, then `Group Configuration`, and then the `iSCSI` tab. The Group Configuration – iSCSI window appears. See Table 8-2.
3. In the iSCSI Authentication panel, select `Enable RADIUS authentication for iSCSI initiators`.
4. Optionally, select `Consult locally defined CHAP accounts first`.

5. If you have not already configured the group to use a RADIUS server, click `RADIUS settings` and add at least one RADIUS server. See the procedure in *Using RADIUS authentication and accounting servers* on page 4-9 for adding RADIUS servers.

6. Click `Save all changes`.

After creating the CHAP account, create an access control record for a volume and specify the CHAP user name in the record. See *Configuring access control records*.

Table 8-2: iSCSI Authentication Panel – RADIUS Authentication Fields

Field	Description	Shortcut	User Action
Enable RADIUS authentication for iSCSI initiators	Enables RADIUS authentication for iSCSI initiators.	Alt+E	None
Consult locally defined CHAP accounts first	Consults locally defined CHAP accounts before using RADIUS authentication.	Alt+C	*Creating a local CHAP account* on page 8-2
RADIUS settings	Launches the RADIUS settings dialog, which specifies RADIUS authentication and accounting servers.	Alt+D	*Modifying RADIUS server settings* on page 4-10

If you want to enable target authentication (for mutual authentication), see *Configuring target authentication* on page 8-4.

Configuring 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, click `Group`, then `Group Configuration`, and then the `iSCSI` tab. The Group Configuration – iSCSI window appears. See the online help for information about the data fields and options.

To change the target authentication user name or password:

1. Click `Modify` and change the user name or password.

2. Enter the target authentication user name and password from Step 2 in the iSCSI initiator configuration interface, where you enable mutual authentication.

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 records to determine if access should be authorized. However, access control records do not prevent multiple initiators, either on the same computer or different computers, from accessing the same target.

Therefore, by default, the group disables multi-host (shared) access to a target. Therefore, 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 4.0 or higher, the group allows multi-host access to targets.

An iSNS (Internet Storage Name Service) server can facilitate iSCSI initiator discovery of iSCSI targets in a SAN.

Configuring the group to use an iSNS server

You can configure a PS Series group to allow an iSNS (Internet Storage Name Service) server to automatically discover and maintain an up-to-date list of group targets. The iSNS server then provides the target names to initiators configured to use the server.

Note: 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.

Requirement: 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 4.1.4 or a higher version. If a member is running a previous firmware version, iSNS servers can automatically discover all group targets.

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.

To display the iSNS configuration in the group, click `Group`, then `Group Configuration`, and then the `iSCSI` tab. The Group Configuration – iSCSI window appears. See the online help for information about the data fields and options.

To configure an iSNS server:

1. In the iSCSI Discovery panel, under iSNS servers, click `Add`.
2. Specify the IP address for an iSNS server and click `OK`. Use the format *`ip_address`*:*`port`* if using a port other than 3205.

   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 specified. Click the up and down arrows to change the IP address order.
3. In the Group iSCSI window, click `Save all changes` (Control+S).

Modifying an iSNS server

To modify the IP address for an iSNS server, select the address in the iSCSI Discovery panel in the Group iSCSI window and click `Modify`. Change the IP address, click `OK`, and then click `Save all changes` (Control+S).

Deleting an iSNS server

To delete the IP address for an iSNS server, select the address in the iSCSI Discovery panel in the Group iSCSI window and click `delete`. Then, click `Save all changes` (Control+S).

Preventing the 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. This can result in a large number of events logged in the group and is an inefficient 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 only discover targets for which they have the correct access credentials.

Enable the iSCSI discovery filter

1. Click `Group`, then `Group Configuration`, and then the `iSCSI` tab.
2. In the iSCSI Discovery panel, under iSNS discovery filter, select `Prevent unauthorized hosts from discovering targets`.
3. Click `Save all changes` (Control+S).

Disable the iSCSI discovery filter

1. Click `Group`, then `Group Configuration`, and then the `iSCSI` tab.
2. In the iSCSI Discovery panel, under iSNS discovery filter, deselect `Prevent unauthorized hosts from discovering targets`.
3. Click `Save all changes` (Control+S).

Multi-host 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 might result in volume corruption.

When an initiator tries to log in to a target, the group uses access control records to determine if access should be authorized. However, access control records do not prevent multiple initiators, either on the same computer or different computers, from accessing the same target.

Therefore, by default, the group disables multi-host (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 4.0 or higher, the group allows multi-host access to targets.

If you disable multi-host access to a volume, when an initiator tries to log in to the volume:

- If there is no iSCSI initiator connection to the volume, the group uses access control records 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 records to determine whether to authorize access.

However, some environment might need multi-host access to a target. You can enable multi-host 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 records as the primary method of protecting iSCSI targets in a group.

You can enable or disable multi-host access when creating a volume. You can also modify a volume or snapshot and enable or disable multi-host access.

Connecting 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.

Both hardware and software iSCSI initiators are available from a variety of vendors. Install and configure an initiator using the vendor-supplied instructions. See your
PS Series support provider for information related to your iSCSI initiator. Also, read the PS Series *Release Notes* for initiator information.

Note: Access to iSCSI targets is through TCP/IP port 3260 (the standard iSCSI port).

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

The general login overview is:

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 records.

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

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.

Note: In some file systems, volumes and snapshots must have read-write permission even if the file system is read-only. See *Modifying volume permission* and *Modifying snapshot permission*.

9 Basic volume operations

Basic volume operations consist of creating volumes, setting up access controls, and modifying volume attributes. You can also clone a volume to create an exact copy.

About volumes

A computer uses an industry-standard **iSCSI initiator** to access a volume in a group. Most operating systems provide an iSCSI initiator (in the form of an iSCSI driver or a host bus adapter), with a range of price and performance options.

To access storage in a PS Series group, you allocate portions of a storage pool to volumes. Each volume appears on the network as an iSCSI target. Computers with iSCSI initiators can connect to the target, which appears as a regular disk.

When you connect to a volume, it appears like a regular disk drive that you can format by using the normal operating system utilities. As the group configuration changes, volumes continue to be accessible through the same iSCSI targets, and no modifications are necessary.

You assign each volume a size (reported size) and a storage pool. The group automatically load balances volume data across pool members. Optionally, you can reserve snapshot space or replication space for a volume.

For each volume (and snapshot), the group generates an iSCSI target name, which you cannot modify. An iSCSI target name includes a prefix (`iqn.2001-05.com.equallogic`), a string, and the volume name. Initiators use the target name to connect to a volume.

The following is an example of an iSCSI target name for a volume with the name `db3`:

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

Access control records and other mechanisms protect your volumes from unauthorized and uncoordinated access by iSCSI initiators. See *iSCSI target security* on page 8-1.

You can use snapshots to protect volume data from mistakes, viruses, or database corruption. To protect against disasters, you can replicate volume data from one group to another.

About volume types

A PS Series group supports the following volume types:

- Standard volume. The default volume type is a standard volume. There are no restrictions on a standard volume. You can enable (and disable) thin provisioning on a standard volume.
- Template volume. 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. When 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 volume. 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. See *About template volumes and thin clones* on page 10-6.

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.

Displaying the iSCSI target name and alias

To display the iSCSI target name and public alias for a volume, click `Volumes` in the lower-left panel, then expand `Volumes`, then select the volume name, and then click the `Connections` tab.

See the online help for information about the data fields and options.

About volume space allocation

It is important to understand how the group allocates space to volumes. This helps you to size volumes correctly. 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.

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:

- No thin provisioning – The volume reserve is equal to the reported size.
- Thin provisioning – 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. See *About thin provisioning*.

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 volume security and access controls

Online volumes and snapshots are seen on the network as iSCSI targets. It is important to understand how to protect your iSCSI targets from unauthorized and uncoordinated access by iSCSI initiators. See *iSCSI target security*.

Access control records prevent unauthorized computer access to iSCSI targets (volumes or snapshots). A volume and its snapshots share a list of access control records. An access control record 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.

When you create a volume, you can create an access control record for the volume. You can later create additional access control records (up to 16) for a volume. See *Configuring access control records* on page 9-10.

In addition, you can allow or disallow volume access from multiple initiators, depending on your configuration needs. See *Allowing or disallowing multi-host volume access* on page 9-14.

About volume data protection

Dell recommends that you use snapshot and replication functionality to protect volume data.

A snapshot is a point-in-time copy of volume data that can protect against mistakes, viruses, or database corruption. You can recover data from a snapshot by setting it online or by restoring the volume from a snapshot. See *Snapshot management*.

To protect against disasters, you can replicate volume data from one group to another. A replica is a point-in-time copy of volume data that is located on a different group and has no dependencies on the original volume. In the event of a disaster, you can host the volume from the recovery group and later failback to the original group, with minimal disruption to users. See *Volume replication*.

Recommendation: Dell recommends that you protect data by using a robust backup application, in addition to snapshot and replication functionality.

Volume attributes

Table 9-1 describes the attributes that allocate space and set the characteristics of a volume. The first column lists volume attributes, and the second column describes them.

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 9-1: Volume Attributes

Volume Attribute	Description
Name	Name, up to 63 alphanumeric characters (including periods, hyphens, and colons), that identifies the volume for administrative purposes. A volume name must be unique in a 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. See *Modifying a volume name or description* on page 9-12.
Description	Optional description for the volume - up to 127 ASCII characters. See *Modifying a volume name or description* on page 9-12
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. See *Moving a volume to a pool* on page 7-5.
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 15. You cannot change the reported size of a template volume. See *About volume space allocation* on page 9-2.

Table 9-1: Volume Attributes (Continued)

Volume Attribute	Description
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 group-wide volume settings. See *About thin provisioning* on page 10-1.
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 group-wide snapshot reserve setting. See *Snapshot management*.
iSCSI alias	Name that some iSCSI initiators display. Administrators use the alias to identify the iSCSI target. The default is the group-wide volume setting. See *Modifying a volume alias* on page 9-13.
Permission	Whether the volume is read-write (the default) or read-only. You cannot set a template volume to read-write permission. See *Modifying volume permission* on page 9-14.
Administrative status	Whether the volume is online (the default) or offline. Initiators cannot discover or connect to an offline volume. See *Setting a volume offline or online* on page 9-13.
Access controls	Conditions that computers must meet to access the volume and its snapshots. To allow volume or snapshot access, you must create at least one access control record. You can do this when creating a volume or after you create the volume. See *About iSCSI target access controls* on page 8-1.
Administrator	You can assign a volume to a specific volume administrator. See *Modifying the administrator for a volume* on page 9-13.
Multi-host access setting	Whether the volume allows or disallows (default) access from initiators with different IQNs. See *Allowing or disallowing multi-host volume access* on page 9-14.
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. See *Enabling or disabling iSNS discovery* on page 9-14.
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. You can override automatic performance load balancing by enabling a RAID level preference on a volume. Thin clones inherit the RAID preference of the template volume. See *Enabling and disabling a volume RAID preference* on page 10-16.
Member binding	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. Thin clones inherit the member binding of the template volume. See *Binding and unbinding a volume to a member* on page 10-16.

Displaying group-wide default volume settings

When you create a volume or enable thin provisioning on a volume, group-wide defaults are applied, unless you explicitly override them for a volume. These default values control snapshot space, snapshot behavior, thin provisioning space, and iSCSI alias naming.

To display group-wide default volume settings, click `Group`, then `Group Configuration`, and then the `Defaults` tab. The Group Configuration – Defaults window appears, displaying the information about the default volume settings. You can change the default volume settings. See *Modifying group-wide volume settings*.

The Group Configuration – Defaults window provides the following information:

- Space allocation percentages.
- Space management policies for snapshot retention and placing the volume offline.
- iSCSI identifiers and volume name use.

You can change a setting for an existing volume. See the online help for information about the data fields and options. See also *Volume attributes* on page 9-3.

Modifying group-wide 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, and iSCSI alias naming. See *Displaying group-wide default volume settings* on page 9-5.

You can modify the group-wide default values to meet the needs of your configuration.

Note: Changes to the group-wide default values apply only to new volumes.

To modify group-wide volume settings:

1. Click `Group`, then `Group Configuration`, and then the `Defaults` tab.
2. In the Group Defaults window, change the default settings.
3. Click `Save all changes` (Control+S).

Creating standard volumes

To provide storage space to end users, you create standard volumes that users can then access from host computers.

Before you create a standard volume, you need to understand:

- Volume attributes and group-wide default settings the group apply to a volume. See *Volume attributes* on page 9-3 and *Displaying group-wide default volume settings* on page 9-5.
- Volume security and access controls. See *iSCSI target security*.
- Risks and benefits associated with thin provisioning before applying this functionality to a volume. See *About thin provisioning*.

To create a standard volume:

1. Click `Volumes` in the lower-left panel and then click `Create volume`.
2. In the Create Volume – Volume Settings dialog box, specify:
   - Volume name
   - Optional volume description
   - Storage pool
3. Click `Next`.
4. In the Create Volume – Space dialog box, specify:
   - Reported volume size and the unit of measure.
   - Whether to use thin provisioning on the volume. See *About thin provisioning*.

     The values in the Pool Space table depend on the reported size and thin provisioning setting. If creating the volume exceeds the capacity of the pool, the table cell showing free pool space displays a negative value. If you enable thin provisioning, the Create Volume – Space dialog box allows you to adjust the default settings for the following thin-provisioned volume attributes, which are based on a percentage of the reported size:

     – Minimum volume reserve

     – In-use space warning threshold

     – Maximum in-use space
   - Snapshot reserve (optional), based on a percentage of the volume reserve.
5. Click `Next`.
6. In the Create Volume – iSCSI Access dialog box, specify:
   - Conditions a computer must match to connect to the volume and its snapshots. Specify a CHAP user name, IP address, or iSCSI initiator name. This information generates an access control record that applies to the volume and its snapshots. See *About iSCSI target access controls* on page 8-1.
   - Permission for the volume, either read-write or read-only.
   - Whether to allow or disallow (default) access to the volume and its snapshots by initiators with different iSCSI qualified names (IQNs). See *Multi-host access to targets* on page 8-6.
7. In the Create Volume – Summary dialog box, review the volume configuration. If the configuration is correct, click `Finish`. Click `Back` to make changes.

To display the new volume, expand `Volumes` in the far-left panel and select the name of the new volume.

Displaying volumes

To display all the volumes in a group, click `Volumes` in the lower-left panel and then `Volumes` in the far-left tree.

the Volumes Summary window provides the following volume information:

- Identifiers – name, storage pool location and accessibility status.
- Capacity – Reported size and reserve space allocation.
- Parameters – Replication partner, snapshot count, connection count and administration account.

See the online help for information about the data fields and options.

You can modify the display:

1. To display the thin clones attached to each template volume, in the far-left column, pull down the `Tree view options` menu and select `Show thin clones under templates`.
2. Select `Organize by volume type` to display template volumes and thin clones in separate categories.

Table 9-2 shows the icons that identify a volume in the Group Manager GUI, depending on the volume's configuration.

Table 9-2: Volume Icons

Volume Icon	Recovery Mode	Description
		Standard iSCSI volume.
		Volume that is set offline, unavailable for operations.
		Volume collection, a set of related volumes.
		Thin provisioned volume.
		Thin clone volume.
		Template volume.

Move the pointer over a volume to display a context message showing the requested volume status and the current volume status.

See *Monitoring volumes and snapshots* on page 15-25.

Displaying volume details

You can display details about a volume such as its status, space usage, and access records.

Displaying volume status

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Status` tab. The Volume Status window appears, containing the following panels:

- General volume information panel – Provides information about volume attributes:

  - Name, status, and security (access)
  - Capacity and reserve
  - iSCSI connections
  - Storage pool and (if a thin clone) template volume
  - Replication identifiers
- Volume and snapshot space panel – Provides information about volume use:
  - Size, capacity use, and reserve allocation
  - Whether the volume is thin provisioned
  - Snapshot reserve and space use
  - Policies for recovery, RAID, load balancing, and member striping
  - Whether the volume is a thin clone, and the amount of shared space

See the online help for information about the data fields and options.

Displaying access control records

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Access` tab. The Volume Access window appears, containing the following information for volumes and snapshots:

- CHAP user name
- IP address or range
- iSCSI initiator name

An asterisk in the CHAP user, IP address, or iSCSI initiator column means that any value meets the condition. If a record shows asterisks in all columns, the volume or snapshot has unrestricted access, which Dell does not recommend.

See the online help for information about the data fields and options. See also *iSCSI target security* on page 8-1.

Displaying volume snapshots

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Snapshots` tab. The Volume Snapshots window appears, containing the following panels:

- Snapshot summary panel – Provides information about:
  - Reserve space and capacity allocation
  - Warning settings and recovery policy
  - Snapshot schedules and schedule status
- Snapshots panel – Provides information about:
  - Creation timestamp, size, and status

  - Collections and schedules that include this snapshot
  - Accessibility (security) and current connections

See the online help for information about the data fields and options.

Displaying volume replication

Displays the replication configuration and replicas for the volume.

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Replication` tab. The Volume Replication window appears, containing the following panels:

- Replication summary panel – Provides information about:
  - Replication partner and failback parameters
  - Replication reserve allocation
  - Transfer status and next replication time
  - Schedule and schedule status
- Remote Replicas panel – You can switch between two views information:
  - Volume replicas (Alt+V) – Provides information about the volume replicas and replica status
  - Replication history (Alt+R) – Provides information about past replication events, including the timestamp, duration, data transferred and transfer parameters.

See the online help for information about the data fields and options. See also *Volume replication* on page 12-1.

Displaying volume collections

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Collections` tab. The Volume Collections window appears containing the following information:

- Collection name, included volumes and storage pool
- Capacity and space use
- Volume status and number of snapshots
- Replication partner

Note: Because template volumes are not supported in volume collections, the `Collections` tab is not applicable to template volumes.

See the online help for information about the data fields and options.

Displaying volume schedules

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Schedules` tab. The Volume Schedules window appears, containing the following panels:

- Schedules summary panel – Provides information about the snapshot and replication schedules for a volume, schedule status and the next scheduled event.
- Snapshot and Replication schedules panel – Provides information about individual schedules, event times and schedule status.

Select `Also show schedules for collections that include the volume` or enter Alt+A to include the snapshot and replication schedules for volume collections that include the volume.

See the online help for information about the data fields and options.

Displaying volume connections

To display iSCSI connections to a volume, click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Connections` tab. The Volume Connections window appears, containing the following panels:

- Volume iSCSI settings panel – Provides information about the iSCSI target name and alias name.
- iSCSI connections panel – Provides information about the iSCSI initiator name and connection statistics.

See the online help for information about the data fields and options.

Displaying thin clones attached to a volume

To display the thin clones attached to a template volume, click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Thin Clones` tab. The Volume Thin Clones window appears, containing the following information:

- Number of thin clones of the template volume
- Number of demoted thin clones of the template volume
- Combined saved space from using thin clones of the template volume

Configuring access control records

Access control records a prevent unauthorized computer access to iSCSI targets (volumes or snapshots). See *About iSCSI target access controls*.

You can set up an access control record when you create a volume. You can later set up additional access control records (up to 16) and apply them to the volume or its snapshots.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Access` tab.

   The Access Control List panel displays the access control records, the conditions in each record, and whether the record applies to the volume or its snapshots.

   Note: An asterisk in the CHAP user, IP address, or iSCSI initiator column means that any value meets the condition. If a record shows asterisks in all columns, the volume or snapshot has unrestricted access, which Dell does not recommend.

2. In the Access Control List panel, click `Add`. The Add Access Control Record dialog box appears.
3. Select the conditions that a computer must meet and specify the required information (CHAP user name, IP address, iSCSI initiator name).
4. Select whether the access control record applies to the volume, its snapshots, or the volume and its snapshots (default).
5. Click `OK`.

Modifying or deleting an access control record

To modify an access control record:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Access` tab.
2. In the Volume Access window, select the record and click `Modify`. Change the information as needed and click `OK`.

To delete an access control record:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Access` tab.
2. In the Volume Access window, select the record and click `Delete`. Confirm that you want to delete the record.

Notes: You cannot change the iSCSI initiator name in an access control record. Instead, you must delete the record and recreate it using the new initiator name.

If you modify or delete a record, computers that met the original conditions might not meet the new conditions and, therefore, might not be able to log in to the target.

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, contents as the original volume at the time of the cloning.

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.

Restriction: If the original volume is a recovery volume, the new volume is not a recovery volume.

See *Volume attributes* for a description of the attributes that apply to a new volume. See *iSCSI target security* for information about volume security.

Cloning a volume

To clone a volume, including a template volume or a thin clone volume:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name.

2. Click `Clone`. The following dialog box appears: Clone Volume – Settings. Specify the new volume name and description.

3. Click `Next`.The following dialog box appears: Clone Volume – Space. Enable or disable thin provisioning (only applicable to standard volumes). You can also change the thin provisioning space allocation settings for the new volume:

   - Minimum volume reserve
   - In-use space warning threshold
   - Maximum in-use space (maximum volume reserve)

   The group updates the values in the Pool Space table, based on the space settings. If the new volume exceeds the capacity of the pool, the table cell showing free pool space becomes red, displays a negative number, and an error message appears.

   Optionally, specify the amount of space, as a percentage of the volume reserve, to reserve for snapshots of the new volume.

4. Click `Next`. The following dialog box appears: Create Volume – iSCSI Access. Specify the following:

   - Conditions that a computer must match to connect to the volume and its snapshots. Specify a CHAP user name, IP address, or iSCSI initiator name. This information generates an access control record that applies to the volume and its snapshots. See *About iSCSI target access controls*.
   - Permission for the volume, either read-write or read-only.
   - Whether to allow or disallow (default) access to the volume and its snapshots by initiators with different iSCSI qualified names (IQNs). See *Multi-host access to targets*.

5. Click `Next`. The following dialog box appears: Clone Volume – Summary. Review the volume configuration.

6. If the configuration is correct, click `Finish`. To make changes, click `Back`.

Click `Volumes` in the far-left panel. The new volume appears in the list of volumes. See *Displaying volumes* on page 9-6.

Note: Enabling or disabling thin provisioning on a volume is an advanced operation. See *Advanced volume operations*.

Modifying a volume name or description

Be aware of the potential impact of modifying a volume name:

- 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 by the original volume name.

Requirement: A volume name must be unique name and can be up to 63 alphanumeric characters (including periods, hyphens, and colons). A volume description can be up to 127 alphanumeric characters.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, then click the `Modify settings`, and then click the `General` tab.
2. In the Modify volume settings – General dialog box, modify the name or the description. The volume name must be unique in the group and can contain up to 63 alphanumeric characters, including periods, hyphens, and colons. The description can contain up to 127 ASCII characters.
3. Click `OK`.

Modifying 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.

When you create a volume, it has an alias only if the group-wide default is to use the volume name as the alias. Otherwise, the volume does not have an alias.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, then click `Modify settings`, and then click the `Advanced` tab.
2. In the Modify volume settings – Advanced dialog box, specify a volume alias in the `Public alias` field. Like its name, the volume alias can contain up to 63 alphanumeric characters, including periods, hyphens, and colons.
3. Click `OK`.

Modifying the administrator for a volume

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, then click the `Modify settings`, and then click the `General` tab.
2. In the Modify volume settings – General dialog box, change the name of the administrator or select `none` in the `Volume administrator` field.
3. Click `OK`.

Setting a volume offline or online

By default, when you create a volume, the group sets the volume online. An iSCSI initiator can only discover or connect to an online volume. To make a volume inaccessible to iSCSI initiators, set the volume offline; the group closes all current iSCSI connections to the volume.

Requirement: To set a volume online, each member that contains volume data must be online.

To set a volume online, click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Set volume online.`

To set a volume offline:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Set volume offline`.
2. Confirm that you want to set the volume offline.

Modifying volume permission

A volume can have read-write or read-only permission, unless it is a template volume.

Requirement: To change a volume permission to read-only, you must first set the volume offline.

Restriction: You cannot set a template volume to read-write permission.

To modify a volume permission:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Set access type`.
2. In the Set Access Type dialog box, change the permission.
3. Click `OK`.

Allowing or disallowing multi-host 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 might result in volume corruption.

You can allow or disallow multi-host (shared) access to a volume. If you disallow multi-host 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 multi-host access to a volume. See *Multi-host access to targets*.

Requirement: Before disallowing multi-host 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 multi-host access.

To allow or disallow multi-host access to a volume:

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Set access type`.
2. In the Set Access Type dialog box, allow or disallow multi-host access.
3. Click `OK`.

Enabling or disabling 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. See *Configuring the group to use an iSNS server*.

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 4.1.4 or a higher version. If a member is running a previous firmware version (for example, if you downgrade a member from V4.1.4), 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 formats:

```
volume select volume_name isns-discovery enable | disable

volume select volume_name snapshot select snapshot_name isns-discovery enable | disable
```

Deleting a volume

You can delete a volume. Space that the group allocated to the volume becomes part of free pool space.

Note: If you delete a volume, the group also deletes its snapshots. However, the group does not delete any volume replicas on the secondary group.

Requirement: You must set a volume offline before you delete it. The group closes any active iSCSI connections to the volume.

Restriction: You cannot delete a template volume if it still has thin clones attached to it or if recovery thin clones exist.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Delete volume`, `Delete template`, or `Delete thin clone`.
2. Confirm that you want to delete the volume and its data.

10 Advanced volume operations

Only knowledgeable group administrators should perform advanced volume operations, including creating thin-provisioned volumes, using template volumes and thin clones, changing a volume size, using volume collections, and scheduling volume operations.

About 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-allocate” or “over-provision” group storage space.

Recommendation: 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 (that is, you can “over-provision” group 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.

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.

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:

- No thin provisioning – The volume reserve is equal to the reported size.

  For example, even if only 10% of a volume is in use, the group allocates the full reported size.

- Thin provisioning – 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% 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 6 GB, the group allocates an additional 10 GB.

– For a volume with a reported size that is less than 100GB, when free volume reserve falls below 6% of the reported volume size, the group allocates an additional 10% of the reported volume size.

Event messages inform you when in-use volume reserve surpasses a user-defined limit and reaches the maximum.

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 group-wide setting is 10%.
- **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 group-wide setting is 60%.

  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 200 GB, when the in-use volume reserve increases by every additional 10 GB.

  – For volumes smaller than 200 GB, when the in-use volume reserve increases by every additional 5%.

  For example, if you create a thin-provisioned volume with a size of 500 GB and set the warning limit to 75%, a warning occurs when the amount of in-use volume reserve is more than or equal to 75% of 500 GB, or 375 GB.

- **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 group-wide setting is 100%.

  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%, 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%, 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 group-wide default volume settings (see *Modifying group-wide volume settings* on page 9-5), override the default values when you create a thin-provisioned volume, or modify a volume and change the settings.

Enabling 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.

Thin provisioning is not appropriate for all environments or volumes. You must fully understand thin provisioning before implementing the functionality on a volume. See *About thin provisioning* on page 10-1.

When enabling thin provisioning on an existing volume, be aware of these issues:

- 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` in the lower-left panel, then expand `Volumes`, then select the volume, then click `Modify settings`, and then click the `Space` tab.
2. In the Modify volume settings – Space dialog box, select `Thin-provisioned volume`. The Pool Space table values change.
3. Optionally, use the sliders to modify the group-wide default thin provisioning space settings:
   - Minimum volume reserve
   - In-use space warning limit
   - Maximum in-use space (maximum volume reserve)

   See *Thin provisioning space settings* on page 10-2.
4. Click `OK`.

Make sure you carefully monitor the space usage for a thin-provisioned volume. See *Monitoring volumes, collections, and snapshots* on page 15-24.

Disabling thin provisioning on a volume

You can disable thin provisioning on a standard volume.

Restriction: You cannot disable thin provisioning on a template volume, thin clone volume, recovery template volume, or a recovery thin clone volume.

Before disabling thin provisioning, consider the following:

- 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.

  This 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% to 500%). If you disable thin provisioning on the volume, the group might decrease the percentage to a more appropriate value, closer to 100%. 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`, then expand `Volumes`, then select the volume, then click `Modify settings`, and then click the `Space` tab.
2. In the Modify volume settings – Space dialog box, de-select `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 cell displays a negative value.
3. Click `OK`.

Modifying the thin provisioning space settings

You can modify the thin provisioning space settings that are described in *Thin provisioning space settings* on page 10-2.

1. Click `Volumes`, then expand `Volumes`, then select the volume, then click `Modify settings`, and then click the `Space` tab.
2. In the Modify volume settings – Space dialog box, use the sliders to adjust the settings for:
   - Minimum volume reserve
   - In-use space warning limit
   - Maximum in-use space (maximum volume reserve)

   The Pool Space table values change, based on the new values. If a change exceeds capacity, the free pool space cell displays a negative value.
3. Click `OK`.

About reported volume size

You can change the reported size of a volume while the volume is online and without disrupting access to the volume.

Warning: 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. 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.

Restriction: 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:

- 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 changes.

  Note: If you are replicating the volume, the secondary group does not recognize the reported size change until the next replication.

Increasing the reported size of a volume

You can increase the reported size of the volume, while the volume remains online.

See *About reported volume size* on page 10-5 for information on the impact of changing the reported size.

1. Click `Volumes`, then expand `Volumes`, then select the volume, then click `Modify settings`, and then click the `Space` tab.
2. In the Modify volume settings – Space dialog box, specify the new reported volume size in the `Volume size` field. If the size you specify is not a multiple if 15MB, the group rounds up the value to the nearest multiple of 15.

   The values in the Pool Space table change, based on the new volume size. If you configured the volume for replication, a table showing delegated space on the replication partner also appears. If the new volume size exceeds the capacity of a pool, the free space cell displays a negative value.
3. For a thin-provisioned volume, optionally modify the in-use warning value and maximum in-use space value using the slider bars. See *Thin provisioning space settings* on page 10-2.
4. Click `OK`.
5. Optionally, confirm that you want to create a snapshot of the current volume prior to the resizing.

Decreasing the reported size of a volume

You can decrease the reported size of the volume, while the volume remains online. Decreasing the size of a volume is sometimes called “shrinking” a volume.

See *About reported volume size* on page 10-5 for information on the impact of changing the reported size of a volume.

Caution: If you decrease a volume size to less than the amount of space currently in use, you can lose data.

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 size
```

See the *CLI Reference* manual for more information about using CLI commands.

About 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 and displays it in the Volume Status window.

See *Space considerations for template volumes and thin clones* on page 10-7.

With a few exceptions, all normal volume operations apply to template volumes and thin clones. See *Restrictions on template volumes and thin clones* on page 10-8.

Thin clones are considered attached to the template volume and cannot exist without it, 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. If you expand `Volumes` in the far-left panel, you can choose to display all volumes in alphabetical order or display thin clones under the template volume on which they depend.

If 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, if you expand `Inbound Replicas` in the far-left panel, you can choose to display thin clone replica sets under the template replica set on which they depend.

In addition, 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.

Because of this dependency, the group does not allow you to 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. Also, 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. Just as with a standard, thin provisioned volume, as you write to a thin clone, the group allocates more space and increases the volume reserve.

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.

In the Volume Status window for a thin clone volume, the Shared Space table (only appears 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 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 10-1.

Table 10-1: 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.
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 there are still recovery thin clone volumes, thin clone replica sets, or permanently promoted thin clone replica sets attached to the volume.

Converting a standard volume to a template volume

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.

Note: 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.

Requirement: Before converting to a template volume, make sure the standard volume contains all the data that is shared with the thin clones. Also, 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 1 GB, and in-use space is 900 MB, you can write approximately 100 MB to each thin clone before you must increase the thin clone size.

1. Click `Volumes`, then expand `Volumes`, then select the volume, and then click `Convert to template`.
2. Confirm that you want to convert to a template volume.

Converting a template volume to a standard volume

Restriction: You cannot convert a template volume to a standard volume if thin clones are attached to the template volume or if the template volume is configured for replication.

Restriction: You must disable replication before you can convert a template volume to a volume.

Note: Space used to store template replicas on the secondary group becomes unmanaged if you convert a template volume to a standard volume.

1. Click `Volumes`, then expand `Volumes`, then select the volume, and then click `Convert to volume`.
2. Confirm that you want to convert to a volume.

Creating a thin clone

1. Click `Volumes`, then expand `Volumes`, then select the template volume, and then click `Create thin clone`.
2. In the Create Thin Clone – Volume Settings dialog box, enter a unique name and optional description and click the `Next` button.
3. In the Create Thin Clone – Space dialog box, change the snapshot reserve setting and the thin provisioning settings, and click the `Next` button.
4. In the Create Thin Clone – iSCSI Access dialog box, create an access control record for the volume, select the permission (read-write or read-only), and specify the multi-host access setting. Then, click the `Next` button.
5. In the Summary dialog box, click the `Finish` button if the thin clone configuration is correct. Click the `Back` button to make changes.

Detaching a thin clone from a template volume

Detaching a thin clone from a template volume breaks the dependency between the thin clone and the template volume.

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. Therefore, when you detach a thin clone, the volume reserve for the thin clone increases by the amount of space the thin clone shares with the template volume.

Note: If 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.

Restriction: You cannot detach a thin clone if replication is enabled for the thin clone.

1. Click `Volumes`, then expand `Volumes`, then select the thin clone, and then click `Detach from template`.
2. Confirm that you want to detach the thin clone.

Displaying template volumes and thin clones

The procedure for displaying information about a template volume or a thin clone is the same as for any volume.

1. Click `Volumes`, then expand `Volumes`, and then click the template volume or thin clone.
2. Click the `Status` tab to display the Volume Status window.
3. Click the other tabs to display additional volume information.

For template volumes, there is no `Collections` tab. You cannot use a template volume in a volume collection.

For template volumes, click the `Thin Clones` tab to display the thin clones that are attached to the template volume.

To display thin clones under the template volume in the far-left panel, pull down the tree view options menu and select `Show thin clones under templates`.

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.

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).

Restriction: You cannot use a template volume in a volume collection.

Creating a volume collection

1. Click `Volumes`, then `Volume Collections`, and then `Create volume collection`.
2. In the Create Volume Collection – General Settings dialog box, specify a name (up to 63 ASCII characters) for the collection and an optional description (up to 127 ASCII characters). Then, click `Next`.
3. In the Create Volume Collection – Components dialog box, select the volumes to include in the collection (up to eight volumes) and click `Next`.

4. Review the Create Volume Collection – Summary dialog box and, if satisfactory, click `Finish`. To make changes, click `Back`.

The volume collection appears in the far-left GUI panel, under `Volume Collections`.

Displaying volume collections

Click `Volumes` and then `Volume Collections`. The Volume Collection Summary window appears, containing the following information:

- Collection, volumes, and storage pool identifiers.
- Capacity and reserve
- Volume status and number of snapshots
- Replication partner

See the online help for information about the data fields and options.

Displaying details for a volume collection

Click `Volumes`, then expand `Volume Collections`, then select the collection. Click the tabs to display details about the volume collection.

Status tab

Click `Volumes`, then expand `Volume Collections`, then select the collection. Click the `Status` tab. The Volume Collection Status window appears, containing the following panels:

- Collection status panel – Provides information about snapshot collections and replication status.
- Collection volumes panel – Provides information about:
  - Volumes in the collection, and the storage pool
  - Capacity and reserve
  - Volume status and number of snapshots
  - Replication partner
  - iSCSI connections

The information is sorted by Volume name. You can sort the information by clicking the column headings.

See the online help for information about the data fields and options.

Snapshots tab

Click `Volumes`, then expand `Volume Collections`, then select the collection. Click the `Snapshots` tab. The Volume Collection Snapshot window appears, containing the following panels:

- Snapshot summary panel – Provides information about snapshot collections, schedules, and scheduled events.
- Snapshots panel – Provides information about snapshot timestamp, schedule name, security, and connections.

See the online help for information about the data fields and options.

Replicas tab

Click `Volumes`, then expand `Volume Collections`, then select the collection. Click the `Replicas` tab. The Volume Collection Replicas window appears, containing the following panels:

- Replication summary panel – Provides information about the replication partner and replication schedules and schedule status.
- Remote replicas panel – Provides information about remote replicas and their status.

See the online help for information about the data fields and options.

Schedules tab

Click `Volumes`, then expand `Volume Collections`, then select the collection. Click the `Schedules` tab

The Volume collection window appears, containing the following panels:

- Schedules summary panel – Provides information about snapshot and replication schedules, schedule status, and scheduled events.
- Snapshot and replication schedules panel – Provides information about objects that a schedule can create and the parameters of the schedule.

You can use the Snapshot and Replication Schedules panel to create, modify, or delete scheduled snapshot and/or replication operations. See the online help for information about the data fields and options.

Modifying a volume collection

1. Click `Volumes`, then expand `Volume Collections`, then select the collection, and then click `Modify volume collection`.
2. In the Modify Volume Collection dialog box:
   - To change the collection name or description, click the `General` tab and modify the name (up to 63 ASCII characters) or description (up to 127 ASCII characters).
   - To add volumes or remove volumes from the collection, click the `Components` tab. Select and deselect volumes.
3. Click `OK`.

Deleting a volume collection

Note: 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`, then expand `Volume Collections`, then select the collection, and then click `Delete volume collection.`
2. Confirm that you want to delete the collection.

Scheduling volume operations

You can create schedules to automatically perform volume 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 or a volume collection.

Note: Using a schedule can generate a large number of snapshots or replicas, so make sure that you have sufficient snapshot or replication space. You can set a limit on the maximum number of snapshots or replicas that a schedule creates.

If a volume is part of a volume collection, make sure a schedule for the collection does not overlap a schedule for the volume.

Restriction: You cannot use manual transfer replication with a replication schedule.

You cannot schedule a snapshot or replication operation for a template volume.

Schedule attributes

Table 10-2 describes the schedule attributes. The first column lists attributes, and the second column describes them. You set the attribute values when you create a schedule. You can later modify the schedule and change the values, if you want.

Table 10-2: Schedule Attributes

Attribute	Description
Name	Identifies the schedule for administrative purposes. Up to 63 ASCII characters. The schedule name must be unique in the group.
Type	Type of schedule, either snapshot or replication.
Enabled or disabled	Whether you want to enable (run) or disable the schedule.
Frequency	How often the schedule runs (once, hourly, or daily). For one-time schedules, you can specify the day and time when the schedule runs. For hourly schedules, you can specify the start and end date, the time when the schedule runs, and how often the schedule runs (from 5 minutes to 12 hours apart). For daily schedules, you can specify the start and end date, how often the schedule runs (from every day to every 100 days), the time when the schedule runs, and how often during the day the schedule runs (from 5 minutes to 12 hours apart).

Table 10-2: Schedule Attributes (Continued)

Attribute	Description
Number of snapshots or replicas to keep	How many snapshots or replicas to keep (from 1 to 512; the default is 10). This attribute applies only to snapshots or replicas that the schedule creates. If the schedule exceeds the maximum number of snapshots or replicas to keep, the group automatically deletes the oldest snapshots or replicas before creating new ones. The group closes any active iSCSI connections to a snapshot before deleting the snapshot. Regardless of the number of snapshots and replicas you choose to keep, the size of the snapshot reserve and the replica reserve limits the number of snapshots and replicas.
Snapshot permission	Applicable only to snapshot schedules, applies either read-write or read-only (the default) permission to the snapshots that the schedule creates.

Creating a schedule

When you create a schedule, you specify the attributes described in Table 10-2.

1. Do one of the following:

   - For a volume, click `Volumes`, then expand `Volumes`, then select the volume, and then click `Create schedule`.
   - For a volume collection, click `Volumes`, then expand `Volume Collections`, then select the collection, and then click `Create schedule.`

2. In the Create Schedule – Schedule Type dialog:

   - Enter a schedule name.
   - Select the type of schedule (snapshot or replication).
   - Select the frequency. Alternately, click `Reuse existing schedule` to use an existing schedule as the basis for the new schedule. You can then modify the schedule attributes.
   - Choose whether you want to enable the schedule.

   Click `Next`.

3. The next dialog box depends on the selected schedule frequency and enables you to control how often the schedule runs, the number of snapshots or replicas to keep, and whether the snapshots are read-write or read-only.

   If you selected `Run once` for a frequency, the Create Schedule – Run Once dialog box appears.

   If you select `Hourly schedule` for a frequency, the Create Schedule – Hourly Schedule dialog box appears.

   If you select `Daily schedule` for a frequency, the Create Schedule – Daily Schedule dialog box appears.

   Then, click `Next`.

4. In the Create Schedule – Summary dialog box, if the schedule configuration is correct, click `Finish`. To make changes, click `Back`.

Displaying volume schedules

1. Click `Volumes`, then expand `Volumes`, then select the volume, and then click the `Schedules` tab.

   The Schedules Summary panel shows the status of the volume schedules and the creation time for the next scheduled snapshot or replica. A running schedule is a schedule that you enabled.

2. In the Snapshot and Replication Schedules panel, select a schedule to display details. To display the schedules for volume collections that include the selected volume, select `Also show schedules for collections that include the volume`.

See the online help for information about the data fields and options.

Modifying a schedule

Table 10-2 describes the schedule attributes. You can modify any attribute except for:

- Schedule type (snapshot or replication)
- Schedule frequency (one time, hourly, or daily)

1. Do one of the following:

   - For a volume, click `Volumes`, then expand `Volumes`, then select the volume, and then click the `Schedules` tab.
   - For a volume collection schedule, click `Volumes`, then expand `Volume Collections`, then select the collection, and then click the `Schedules` tab.

2. Select the schedule in the Snapshot and Replication Schedules panel and click `Modify`.
3. Change the schedule attributes.
4. Click `OK`.

Deleting a schedule

Deleting a schedule does not affect the snapshots or replicas that the schedule created.

1. Do one of the following:

   - For a volume, click `Volumes`, then expand `Volumes`, then select the volume, and then click the `Schedules` tab.
   - For a volume collection schedule, click `Volumes`, then expand `Volume Collections`, then select the collection, and then click the `Schedules` tab.

2. Select the schedule in the Snapshot and Replication Schedules panel and click `Delete`.
3. Confirm that you want to delete the schedule.

Enabling and disabling 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.

You can override automatic performance load balancing by enabling a RAID level preference (RAID 10, RAID 50, RAID 5, or RAID 6) 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.

If you disable a RAID preference on a volume, the group resumes automatic performance load balancing.

Restriction: Thin clones inherit the RAID preference, if any, of the template volume. You cannot set a separate RAID preference for a thin clone.

Requirement: To enable a volume RAID preference, make sure 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.

See *Displaying storage pools* on page 7-3 to display the RAID levels of the pool members.

1. Click `Volumes`, then expand `Volumes`, then select the volume, then click `Modify settings`, and then click the `Advanced tab`.
2. In the Modify volume settings – Advanced dialog box, under Volume RAID preference, select the preferred RAID level or select `Automatic` to disable a RAID preference.
3. Click `OK`.

Binding and unbinding a volume to a member

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.

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.

Restriction: 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 use the Group Manager GUI to bind a volume to a member. Instead, you must use the following Group Manager CLI command format:

```
volume select volume_name bind member_name
```

To unbind a volume from a member, use the following CLI command format:

`volume select` *`volume_name`* `unbind`

See the PS Series *CLI Reference* manual for more information about using CLI commands.

Managing a volume or snapshot with lost blocks

In rare circumstances, a volume (or snapshot) might lose blocks. For example, this can occur if there is a power failure 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.

1. Click `Volumes`, then expand Volumes, then select the volume name, and then click the `status` tab.
2. Click `offline-lost-cached-blocks`.
3. In the dialog box that appears, do one of the following:
   - 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`.

     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.
   - 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`.
4. Click `OK`.

   Note: Setting a volume with lost blocks online is a data integrity risk. The blocks might contain old or invalid data.

11 Snapshot management

Snapshots greatly simplify and increase the performance of backup and recovery operations.

About snapshots

A snapshot is a point-in-time copy of volume data. Creating snapshots on a regular basis can protect you from data loss due to mistakes, viruses, or database corruption.

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.

Creating a snapshot does not prevent access to a volume, and the snapshot is instantly available to authorized iSCSI initiators. 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.

Restriction: You cannot restore a template volume from a snapshot.

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: If a volume is offline, all its snapshots are also offline. If you delete a volume, the group deletes all its snapshots.

Like volumes, snapshots appear on the network as iSCSI targets. All the iSCSI target security mechanisms apply to snapshots. See *iSCSI target security* on page 8-1.

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. See *Restoring a volume from a snapshot* on page 11-9.
- Clone a snapshot. The new volume contains the volume data that existed at the time you created the snapshot. See *Cloning a snapshot to create a new volume* on page 11-9.
- Set the snapshot online. iSCSI Initiators can access the target in the usual way. See *Setting a snapshot online or offline* on page 11-11.

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 s**napshot 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 or snapshot collections at a specific time or time interval, such as hourly, daily or weekly.

About snapshot reserve allocation

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. Depending on the policy that you set for snapshot space recovery, the group preserves snapshot reserve as described in *About snapshot reserve settings*.

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 this happens, 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 records, which prevent unauthorized iSCSI initiator access to a volume and its snapshots. See *iSCSI target security*.

About snapshot reserve settings

There are group-wide default values for the following snapshot reserve settings, 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 group-wide default value. You can modify the snapshot reserve value.
- **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.

  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 might be at risk of deletion. To preserve the data in a snapshot, you can clone the snapshot. See *Cloning a snapshot to create a new volume* on page 11-9.
- **Snapshot space warning percentage** – Percentage of the snapshot reserve, when reached by in-use snapshot reserve, results in an event message. The default is 90% of the snapshot reserve.

  For example, if snapshot reserve space is 200 MB and the warning level is 90%, a warning occurs when in-use snapshot reserve equals or exceeds 180 MB.

See *Modifying snapshot reserve settings for a volume* on page 11-3.

About snapshot schedules

You can set up a schedule for creating snapshots of a volume or volume collection at a specific time or on a regular basis.

Using a schedule can cause a large number of snapshots. Make sure you have sufficient snapshot reserve. You can set a limit on the number of snapshots the schedule can create. In addition, the size of the volume's snapshot reserve limits the number of volume snapshots.

Restriction: You cannot schedule snapshots for a template volume.

See *Scheduling volume operations* for information about creating a snapshot schedule.

Modifying snapshot reserve settings for a volume

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Modify snapshot settings`.
2. In the Modify Snapshot Settings dialog box, modify the snapshot reserve, space recovery policy, or the snapshot space warning percentage.

   If you change the snapshot reserve, the values in the Pool Space table change. If the new snapshot reserve value exceeds the capacity of the pool, the free pool space cell displays a negative value.
3. Click `OK`.

Creating snapshots

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.

To create snapshots of multiple volumes at the same time, see *About snapshot collections* on page 11-5.

Requirement: Before you can create a snapshot, you must allocate snapshot reserve for the volume. See *About snapshot reserve settings* on page 11-2.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click `Create snapshot now`.
2. In the Create Snapshot dialog box:
   - Enter an optional description for the snapshot (up to 127 ASCII characters).
   - Select whether to keep the snapshot offline (default) or set the snapshot online.
   - Select whether to make the snapshot permission read-write (the default) or read-only.
3. 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, `dbase-2009-03-25-15:31:14.7668`). Snapshots appear under a volume in the far-left panel listed by 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* pane.

Displaying snapshots for a volume

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree.
2. Select the volume name, and then click the `Snapshots` tab. The Snapshot Summary panel provides the following information:

- Snapshot reserve and capacity used.
- Warning and recovery policy
- Number of, and schedule for snapshots
- Schedule details and events

3. Move the pointer over a snapshot to display a context message showing its name, the current snapshot status, and the requested snapshot status.
4. See Table 15-23 and Table 15-24 for snapshot status.

See the online help for information about the data fields and options.

Displaying snapshot details for a volume

Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left tree, then select the volume name, and then click the `Snapshots` tab. The Snapshots panel provides the following information:

- Collection and schedule containing this snapshot
- Size and status of the snapshot
- Accessibility of (security) and connections to the snapshot

If the snapshot belongs to a snapshot collection, click `View snapshot collection` to see details of the snapshot collection. This link appears in the General Snapshot Information panel only if applicable.

See the online help for information about the data fields and options.

Displaying details of an individual snapshot

1. Click `Volumes` in the lower-left panel and then expand `Volumes` in the far-left tree.
2. Expand a volume name and then click the snapshot. The Snapshot *name* window contains the Snapshot Status and Snapshot Access tabs.

Snapshot status tab

The Snapshot Status tab provides the following panels:

- General snapshot information panel – Provides information about:
  - Volume status and accessibility (security)
  - Schedule and collection relating to this snapshot
  - Creation timestamp and size of the snapshot
- Snapshot iSCSI settings panel – Provides information about the iSCSI target and alias name of the snapshot.
- iSCSI connections panel – Provides information about the initiator and connection statistics.

See the online help for information about the data fields and options.

Snapshot access control tab

The Snapshot Access tab contains the Access Control List panel, which provides the following information:

- Whether the access rules apply to both volume and snapshots
- The methods used: CHAP, IP, or initiator name.

A volume and its snapshots share a list of access control records. See *About iSCSI target access controls* and *Configuring access control records*.

In the Snapshot Access window, select a record in the Access Control List panel to display details, including whether the record applies to the volume or its snapshots.

See the online help for information about the data fields and options.

About snapshot collections

Creating snapshots of multiple volumes simultaneously is useful when you want to protect data in multiple, related volumes. You can do this in one operation, using one of two methods:

- Create snapshots of all the volumes in a volume collection. The resulting set of snapshots, one for each volume in the collection, is called a *snapshot collection*. You can also schedule snapshot collections.
- Create snapshots of multiple volumes without using a volume collection. The resulting set of snapshots, one for each volume, is called a *custom snapshot collection*. You cannot schedule custom snapshot collections.

Creating a snapshot collection

You can create snapshots of all the volumes in a volume collection in one operation. The resulting set of snapshots, one for each volume in the collection, is called a snapshot collection.

Requirement: Before you create a snapshot collection, you must allocate snapshot reserve for each volume in the volume collection. See *About snapshot reserve settings* on page 11-2.

1. Click `Volumes`, then expand `Volume Collections`, then select the collection name, and then click `Create snapshot now.`
2. In the Create Snapshot Collection dialog box, enter an optional description (up to 127 ASCII characters).
3. Click `OK`.

Snapshot collections appear under the volume collection name in the far-left panel.

The default snapshot collection name is the volume collection name, followed by the date and time when you created the snapshots (for example, `datavols-2009-03-25-15:31:14.7668`). However, when snapshot collections appear under a volume collection, you can identify them only by timestamp.

Displaying snapshot collections

1. Click `Volumes`, then expand `Volume Collections`, then select the collection, and then click the `Snapshots` tab.

   The Snapshot Summary panel shows the number of snapshot collections for the volume collection, the most recent snapshot collection timestamp, and any snapshot schedules for the volume collection.
2. Expand a snapshot collection in the Snapshots panel to display the individual snapshots that comprise the collection. A snapshot that is part of a snapshot collection also appears under its volume name in the far-left panel.
3. Place the pointer over the snapshot collection or a snapshot to display details.

Displaying snapshot collection details

1. Click `Volumes`, then expand `Volume Collections`, then expand the collection.
2. Select the snapshot collection timestamp. The Volume Collection – Snapshot Collection window appears.

The Snapshot Collection Status panel describes the snapshot collection, including the volume collection name, the time you (or a schedule) created the snapshot collection, the original number of snapshots in the collection, and the collection Integrity status and Modification status.

The snapshot collection Integrity status can be:

- `complete` – A snapshot exists for each volume in the collection.
- `incomplete` – A snapshot that was originally part of the collection does not exist.

The snapshot collection Modification status can be:

- `not modified` – No snapshot in the collection is set online with read-write permission.
- `potentially modified` – One or more snapshots in the collection are currently set online with read-write permission.

The Snapshots panel displays information about the snapshots in the collection. Double-click a snapshot to display details about the snapshot.

Creating 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.

Requirement: Before you create a custom snapshot collection, you must allocate snapshot reserve for each volume in the volume collection. See *About snapshot reserve settings* on page 11-2.

1. Click `Volumes` and then `Custom snapshot collections`.
2. In the Create Custom Snapshot Collection – General Settings dialog box, enter a name for the custom snapshot collection (up to 63 ASCII characters) and an optional description (up to 127 ASCII characters). You cannot specify spaces in the name. Then, click `Next`.
3. In the Create Custom Snapshot Collection – Components dialog box, select the volumes and click `Next`.
4. In the Create Custom Snapshot Collection – Summary dialog box, review the information. If correct, click `Finish`. 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-2009-05-14-15:08:18.3`).

To display custom snapshot collections, see *Displaying custom snapshot collections* on page 11-7.

Displaying custom snapshot collections

To see information about a custom snapshot collection, you must know the timestamp for the custom snapshot collection, which is the same as the timestamp for its snapshots.

Note: A custom snapshot collection is not associated with a volume collection.

Click `Volumes` and then expand `Custom Snapshot Collections`. The Custom Snapshot Collections panel provides the following information:

- The name, timestamp, and status of the custom snapshot.
- Accessibility (security) of, and connections to the custom snapshot

See the online help for information about the data fields and options.

Displaying snapshot collection status

Click `Volumes`, then `Custom Snapshot Collections`, and then the timestamp of the collection. The Snapshot Collection *name* window contains the following panels:

- Snapshot collection status panel – Provides information about:
  - Name and timestamp of the snapshot collection and any related volume collection
  - Number of snapshots in the collection
  - The integrity and modification status of the collection
- Snapshots panel – Provides information about:
- Name of the snapshot, the related volume, and the storage pool.
- Size and status of the snapshot.
- Accessibility (security) of, and connections to the snapshot.

Modifying a snapshot collection name or description

Requirement: A snapshot collection name can be up to 63 ASCII characters. A snapshot collection description can be up to 127 ASCII characters.

The snapshot collection name appears in the title of the Volume Collection – Snapshot Collection window. Place the pointer over a snapshot collection in the Snapshots panel to see the description.

1. Do one of the following:
   - If you are modifying a snapshot collection name or description:

     Click `Volumes`, then expand `Volume Collections`, then expand the collection, and then select the timestamp for the snapshot collection.
   - If you are modifying a custom snapshot collection name or description:

     Click `Volumes`, expand `Custom Snapshot Collections`, and then select the timestamp for the custom snapshot collection.
2. Click `Modify snapshot collection`.
3. In the Modify Snapshot Collection Settings dialog box, modify the snapshot collection name or description and click `OK`.

Deleting a snapshot collection

Deleting a snapshot collection also deletes the snapshots in the collection.

1. Do one of the following:
   - If you are deleting a snapshot collection

Click `Volumes`, then expand `Volume Collections`, then expand the collection, and then select the timestamp for the snapshot collection.

- If you are deleting a custom snapshot collection:

  Click `Volumes`, then expand `Custom Snapshot Collections`, and then select the timestamp for the custom snapshot collection.

2. Click `Delete snapshot collection`.
3. Confirm that you want to delete the snapshot collection or the custom snapshot collection.

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.

Note: Before a volume restore operation starts, the group automatically creates a snapshot of the current volume.

Requirement: To restore a volume from a snapshot, all members that contain data from a volume or snapshot must be online.

Restriction: You cannot restore a template volume from a snapshot.

1. Disconnect any iSCSI initiators from the volume. Follow the instructions for your operating system and initiator.
2. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, and then select the volume.
3. Click `Set volume offline` and confirm that you want to set the volume offline.
4. Click `Restore volume`.
5. In the Restore Volume from Snapshot dialog box, select the snapshot and click `OK`.
6. In the Restore Volume Confirmation dialog box, choose whether to set the volume online after the restore operation completes (the default) and confirm that you want to restore the volume.

The restored volume has the same name and iSCSI target name as the original volume.

Cloning a snapshot to create a new volume

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, contents as the original volume at the time you created the snapshot.

The group allocates space equal to the volume reserve 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.

See *Volume attributes* for attributes that apply to a new volume. In addition, you should fully understand volume access controls. See *iSCSI target security*.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then select the volume, then select the snapshot timestamp, and then click `Clone snapshot`.
2. Follow the prompts in the Clone Snapshot wizard. See *Cloning a volume* on page 9-11 for information about the information you specify in the wizard dialog boxes.

Modifying snapshot properties

You can modify the properties of a snapshot, including the snapshot name, description, and public alias (public name).

Modifying a snapshot name or description

The default snapshot name is based on the volume name and the time you created the snapshot. You can modify this name and the optional snapshot description.

Note: If you modify a snapshot name, and the public alias (public name) is set to be the same as the snapshot name (as described in *Displaying group-wide default volume settings* on page 9-5), the alias changes to match the new name. The iSCSI target name for the snapshot does not change if you change the snapshot name.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, then select the snapshot timestamp, and then click `Modify snapshot properties.`
2. Click the `General` tab and enter the new snapshot name or description in the Modify Snapshot Properties – General dialog box.
3. Click `OK`.

Modifying the snapshot alias

To help you identify the snapshot, you can modify its public alias (public name). Some iSCSI initiators show the alias along with the iSCSI target name.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, then select the snapshot timestamp, and then click `Modify snapshot properties`.
2. Click the `iSCSI` tab and enter the new alias in `Public alias` field in the Modify Snapshot Properties – iSCSI dialog box.
3. Click `OK`.

Setting 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 records.

If you set a snapshot offline, any current iSCSI connections to the snapshot are lost.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, and then select the snapshot timestamp.
2. Click `Set snapshot online` or `Set snapshot offline`.
3. Confirm you want to set the snapshot online or offline.

Modifying snapshot permission

A snapshot can have read-only or read-write permission.

If you write to a snapshot, it might no longer represent the contents of the base volume at the time of snapshot creation.

Requirement: To change the permission of an online snapshot to read-only, first set the snapshot offline.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, then select the snapshot timestamp, and then click `Set access type.`
2. In the Select Access Type dialog box, select the permission for the snapshot.
3. Click `OK`.

Allowing or disallowing multi-host 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 multi-host (shared) access to a snapshot. If you disallow multi-host 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 multi-host access to a snapshot. See *Multi-host access to targets*.

Note: To disable multi-host access to a snapshot, first disconnect all initiators from the snapshot except one. If multiple initiators have connections when you try to disable multi-host access, the operation fails, unless the initiators have the same IQN.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, then select the snapshot timestamp, and then click `Set access type`.
2. In the Select Access Type dialog box, allow or disallow multi-host access.
3. Click `OK`.

Deleting snapshots

Note: If you delete a snapshot that is part of a snapshot collection or a custom snapshot collection, the collection Integrity status changes to `incomplete`.

1. Click `Volumes` in the lower-left panel, then expand `Volumes` in the far-left panel, then expand the volume name, then select the snapshot timestamp, and then click `Delete snapshot`.
2. Confirm that you want to delete the snapshot.

12 Volume replication

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. When the original volume becomes available, you can failback to the original group.

About replication

An effective data recovery solution must help you correct day-to-day mistakes (such as when users erroneously delete files or volumes), computer viruses, and site disasters.

Some solutions are time-consuming and involve backing up data and manually transporting the backups to a different physical location. Other solutions rely on expensive hardware and the ability to copy data over long distances, which can decrease application performance.

Using the replication technology provided by PS Series firmware, you can 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.

You can use PS Series replication functionality alone or in conjunction with Auto-Snapshot Manager (for Windows or VMware) or Storage Replication Adaptor for VMware Site Recovery Manager. See the product documentation for details.

About replicas

Similar to a snapshot, a **replica** represents the contents of the 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 at the current 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.

How replication works

Before you can replicate volume data, you must configure the group where the volume resides and the group that stores the volume replicas as **replication partners**.

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.

When 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 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.

About manual transfer replication

If you are transferring a large amount of data and your network link between the primary and secondary groups is not sufficient, you can use **manual transfer replication** for a replication operation. Manual transfer replication requires manual tasks and uses external media to copy data to the secondary group, instead of using the network.

For example, the first replication of a volume copies the contents of the volume to the secondary group. If in-use volume space is large and the network is slow, the replication operation can take a long time. In this case, you might want to use manual transfer replication.

You also might want to use manual transfer replication after you defragment a large volume.

Manual transfer replication should only be used when necessary. A properly constructed and sized network should be able to handle network replication.

Manual Transfer Utility

To use manual transfer replication and transfer data between replication partners by using external media, you must first download and install the Manual Transfer Utility from the EqualLogic customer support site. For more information, see the Manual Transfer Utility Installation and User's Guide.

To start the manual transfer wizard:

1. Click `Tools` to display the tools menu.
2. Click `Manual transfer utility`.

The main window appears, containing the following panels:

- Manual Replication in Current Group panel – Provides information about the manual transfer replication operations for the group to which the computer is currently connected:
  - Volume name and direction of transfer
  - Replica timestamp and status
  - Pending actions
- Data Transfers on Local Machine panel – Provides information about the manual transfer operations in the group that are running on the local computer:
  - Group and volume name
  - Current operation, its status and progress
  - Available user actions

Click the checkbox next to `Show data transfers launched from other group` to display all manual transfer operations that are running on the local computer

See the online help for information about the data fields and options.

Replication configuration options

A group can have multiple replication partners. However, you can replicate a volume only to one replication partner at a time. Choose the replication configuration that is right for your environment.

Replication to one partner

One replication partner replicates volumes to another partner. For example, in Figure 12-1, GroupA replicates Volume1 and Volume2 to GroupB. GroupA is the primary group, and GroupB is the secondary group.

Figure 12-1: Replication to One Partner

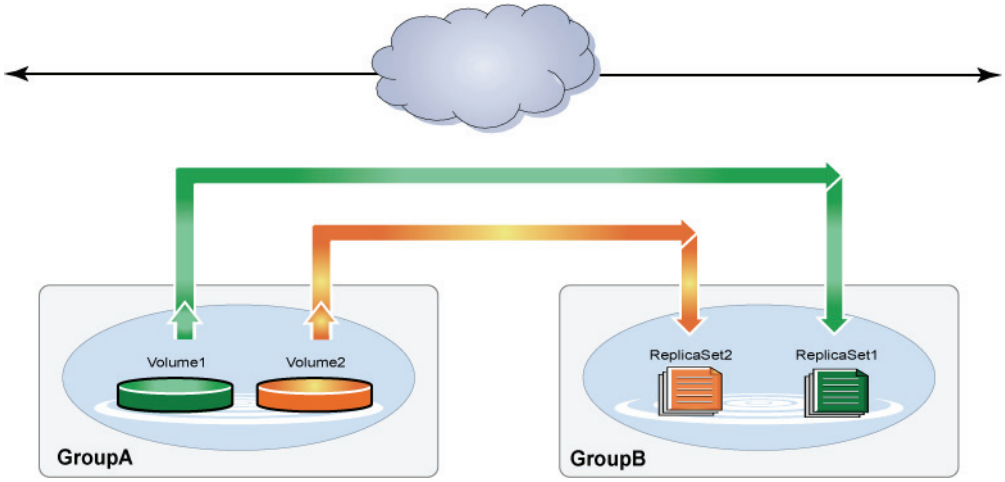

Replication to multiple partners

One replication partner replicates different volumes to different partners. For example, in Figure 12-2, GroupA replicates Volume1 to GroupC, and GroupA replicates Volume2 to GroupB. GroupA is the primary group, and GroupB and GroupC are secondary groups.

Figure 12-2: Replication to Multiple Partners

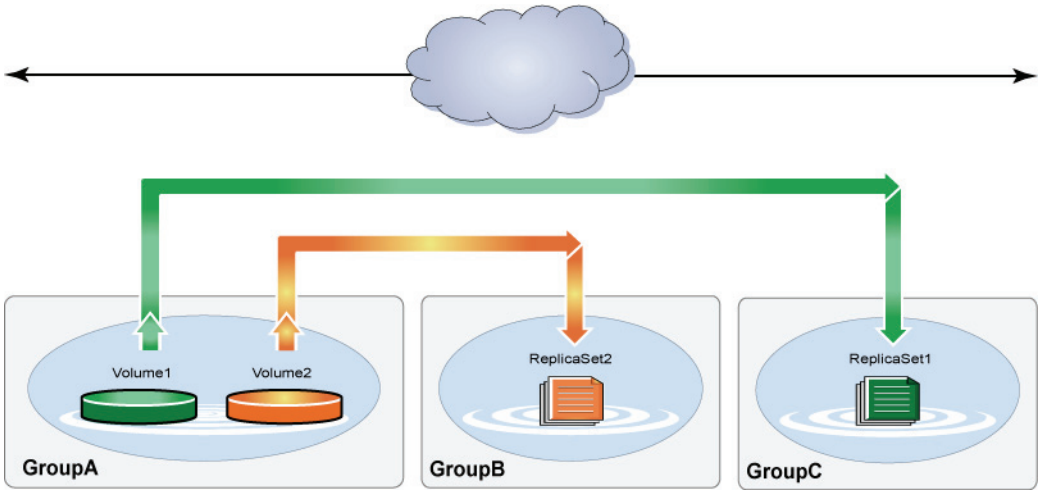

Reciprocal replication between partners

Both partners replicate volumes to each other. For example, in Figure 12-3, 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 12-3: Reciprocal Replication Between Partners

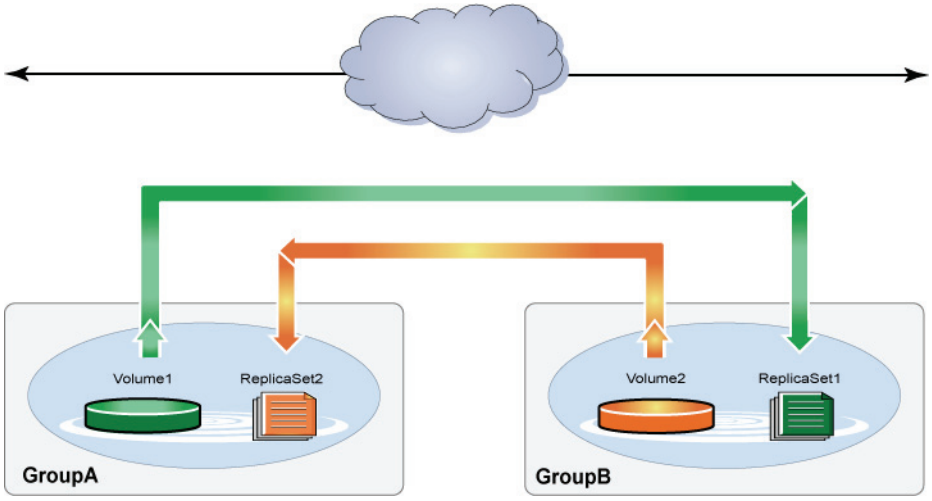

Centralized replication

Multiple partners replicate volumes to another partner. For example, in Figure 12-4, GroupA and GroupB replicate volumes to Group C. In this configuration, GroupA and GroupB are primary groups, and GroupC is the secondary group.

Figure 12-4: Centralized Replication

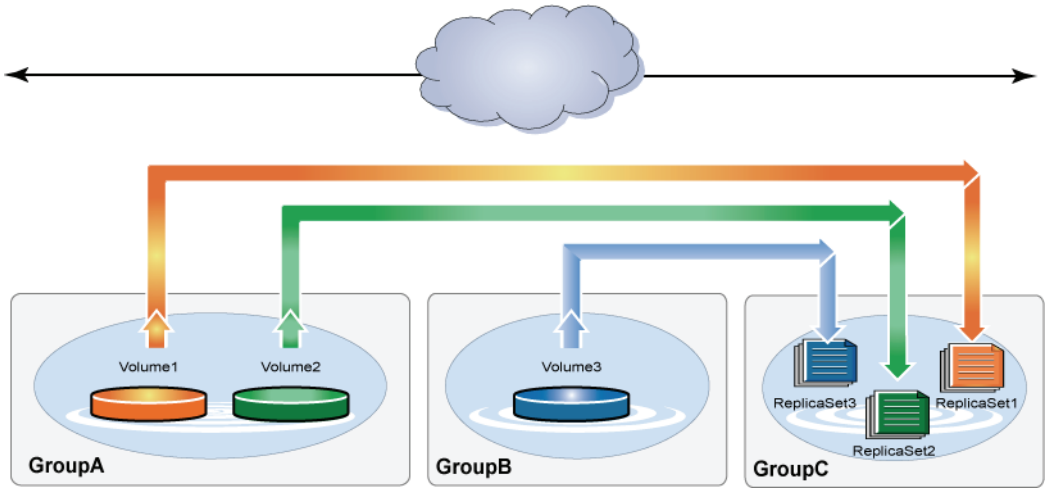

How volume changes affect replication space

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. See *About manual transfer replication* on page 12-3.

For each replication operation, you can display the amount of data that the primary group is transferring to a replication partner. See *Displaying replication activity and replicas for a volume* on page 12-28. You can also display the replication history for a volume and the amount of data transferred for each replication operation.

Best practice for replicating volumes

To help ensure successful replication, for each volume that you want to replicate, follow these steps to set up your replication environment:

1. Plan the volume replication configuration:
    a. Gather the following information to help you determine how much replication space you need:
        - Number of replicas you want to keep and the average time span between each consecutive replica
        - Reported size of the volume
        - Whether thin-provisioned
        - Estimated rate of volume changes (depends on volume usage)
    b. Make sure that the primary group has enough free pool space for the local replication reserve for each replicated volume. See *About local replication reserve* on page 12-8.
    c. Identify a replication partner (secondary group) to store the volume replicas. This secondary group must meet the space and network connectivity requirements in *Replication partner requirements* on page 12-16.

See the PS Series *Release Notes* for replication limits.

2. If you did not already configure the groups as replication partners:

   a. Log in to the primary group and configure the secondary group as a replication partner.

   b. Log in to the secondary group and configure the primary group as a replication partner. Make sure you delegate the correct amount of space to the primary group for storing replicas of primary group volumes.

   If you already configured the groups as replication partners, check the secondary group space that is delegated to the primary group and increase it, if necessary.

   See *About replication partners* on page 12-16.

3. On the primary group, configure the volume for replication, specifying the appropriate replication space values. See *Configuring a volume for replication* on page 12-25.

4. On the primary group, replicate the volume. See *Creating a replica* on page 12-28.

   You can set up a schedule to create replicas on a regular basis. See *Using schedules to create replicas* on page 12-30.

5. Monitor each replication operation and make sure it is successful. See *Displaying replication activity and replicas for a volume* on page 12-28.

   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.

6. 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.

To recover volume data from replicas, see *Data recovery* on page 13-1.

About replication space

Volume replication between partners requires space on the primary group (volume location) and on the secondary group (replica location):

- **Local replication reserve.** Each volume requires primary group space for use during replication and, optionally, for storing the failback snapshot. See *About local replication reserve* on page 12-8.
- **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.

See *About delegated space and replica reserve* on page 12-11 and *Replica reserve usage* on page 12-12.

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. This helps 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 12-5.

Figure 12-5: Local Replication Reserve

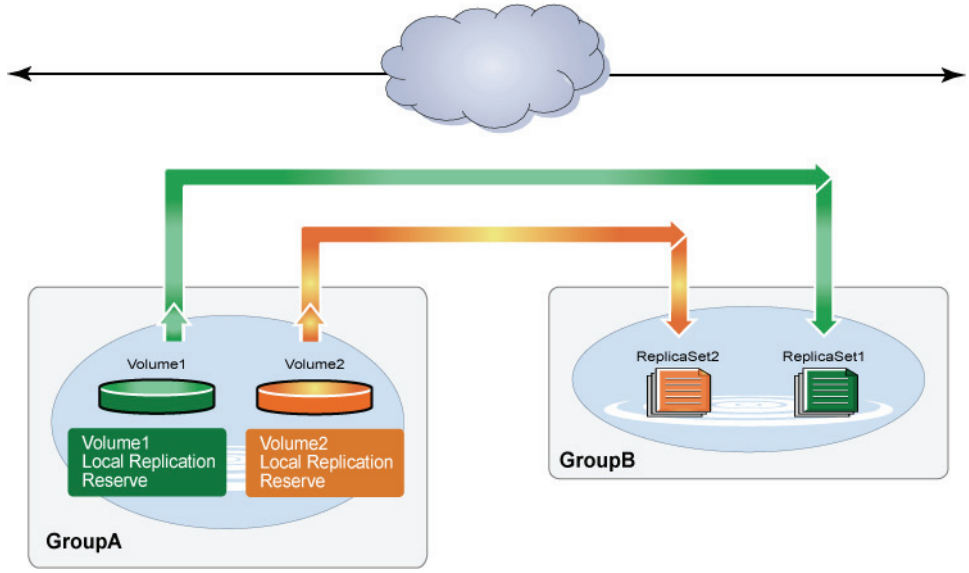

Local replication reserve has two purposes:

- **Preserve the contents of the volume at the time replication started**. If volume writes occur during a replication operation, 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. As volume changes occur during replication, the snapshot consumes 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 failover 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 there is not enough free local replication reserve to complete a replication operation, one of the following occurs:

- If you enabled the option to borrow free pool space, and sufficient free pool space is available (at least 10% free pool space), 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 there is not enough free pool space, the primary group cancels the replication and generates an event message, stating that the replication was cancelled.

If there is not enough free local replication reserve to maintain the failback snapshot, one of the following 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 there is not enough free pool space, 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 there is insufficient local replication reserve to store the changes, the primary group generates an event message advising you to increase the space.

Guidelines for sizing local replication reserve

On the primary group, you 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 there is not enough local replication reserve. You can later modify these settings.

The local replication reserve size is based on a percentage (5% to 200%) 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% for the local replication reserve.
- Keep the failback snapshot. If you want to keep the failback snapshot, specify 200% for the local replication reserve.

However, 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. See *How volume changes affect replication space* on page 12-6.

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.

See *Sizing local replication reserve for use during replication* and *Sizing the local replication reserve for the failback snapshot*.

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% 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%.

  To obtain an appropriate value, estimate the average volume changes that occur during a typical replication operation. Then, use this equation, where 5% is the minimum local replication reserve:

  ```
  5% + change rate
  ```

  For example, if you estimate that at most 10% of the volume changes, a value of 15% might be appropriate (5% plus 10%).

  If you use a local replication reserve value that is less than 100%, 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 there is not enough local replication reserve to complete a replication.

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% ensures sufficient local replication reserve even if the entire volume changes between replication operations.
- **Space-efficient value**. If few volume changes are expected on average between replication operations, use a value less than 100%,

  To obtain an appropriate value, estimate the average volume changes that occur between consecutive replication operations. Then, use this equation, where 5% is the minimum local replication reserve:

```
5% + change rate
```

For example, if you estimate that at most 20% of the volume changes, a value of 25% might be appropriate (5% plus 20%).

If you use a local replication reserve value that is less than 100%, 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 there is not enough local replication reserve to maintain the failback snapshot.

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 Group B volumes to GroupC, GroupC must delegate space to GroupA and GroupB. See Figure 12-6.

Figure 12-6: Delegated Space

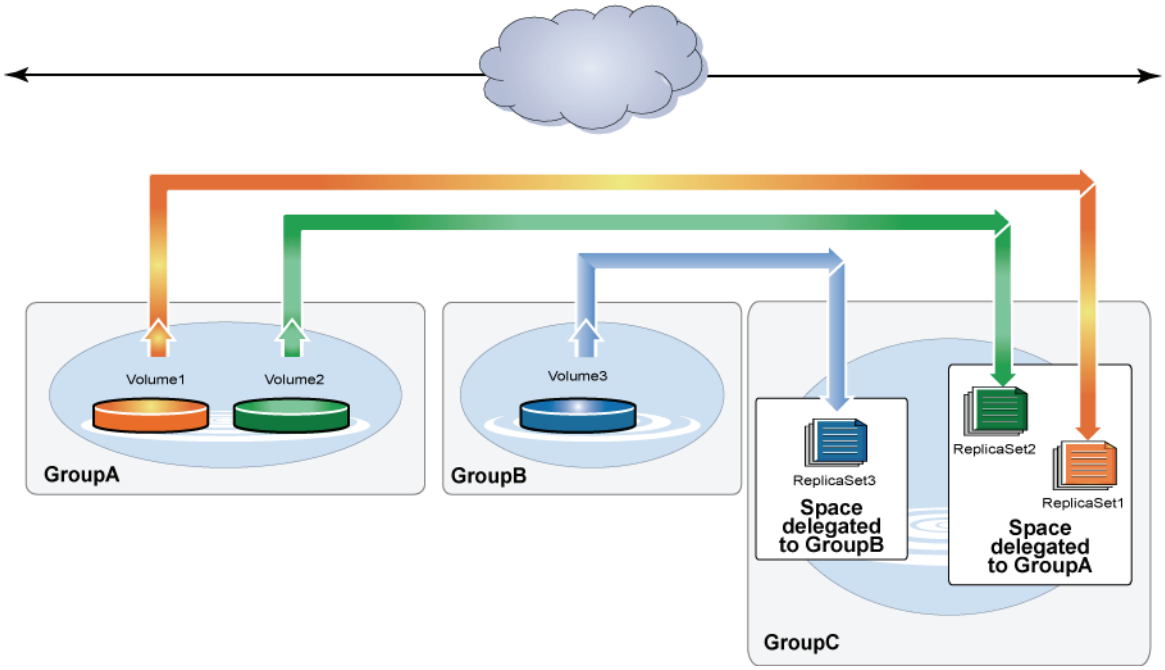

The secondary group administrator delegates space to the primary group when configuring the group as a replication partner. The administrator can modify the partner configuration and increase or decrease delegated space.

When the primary group administrator configures a volume for replication, the administrator assigns a portion of delegated space to the volume. This space, called **replica reserve**, limits the number of replicas that you can keep on the secondary group. You can modify the volume replication configuration and increase or decrease the replica reserve value.

To determine the correct 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. Replica reserve is based on a percentage of the volume's replica volume reserve. See *Replica volume reserve* and *Replica reserve usage*.

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% of the reported volume size.

  For example, if you have a 10 GB volume that is not thin-provisioned, the initial replica volume reserve is 1 GB (10% of 10 GB).

- For thin-provisioned volumes, the current volume reserve.

  For example, if you have a 10 GB volume that is thin-provisioned with a volume reserve of 2.5 GB, the initial replica volume reserve is 2.5 GB.

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, the replica volume reserve increases by the amount of new data added to the volume since the start of the previous successful replication, up to the reported volume size. That is, the replica volume reserve increases by the amount of data that must be replicated.

  For example, if you have a 10 GB volume that is not thin-provisioned, and initiators write 2 GB to the volume, the replica volume reserve increases by 2 GB.

- For thin-provisioned volumes, the replica volume reserve increases to the current size of the volume reserve, up to the maximum in-use space value.

  For example, if you have a 10 GB volume that is thin-provisioned, and the volume reserve increases from 2.5 GB to 4.0 GB, the replica volume reserve increases by 1.5 GB.

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. See *Replica reserve usage*.

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%) of the replica volume reserve, which approximates in-use volume space. As volume usage increases, the replica volume reserve increases; therefore, the replica reserve also increases, providing more space for replicas, up to a limit. See *Replica volume reserve*.

For example, if you specify 200% for the replica reserve, and the replica volume reserve is 2.5 GB, the replica reserve size is 5.0 GB. If the replica volume reserve increases to 6.0 GB, the replica reserve size increases to 12.0 GB.

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. See *Replica reserve usage – first replication* and *Replica reserve usage – subsequent replications*.

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 reserve increased since you enabled replication on the volume.

   Note: If there is not enough free delegated space for the replica reserve increase, the primary group generates an event message, and the replication pauses. Replication continues automatically once there is sufficient delegated space.

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 10 GB of data, free replica reserve decreases by 10 GB.

   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 reserve increased since the previous replication operation.

   Note: If there is not enough free delegated space for the replica reserve increase, the primary group generates an event message, and the replication pauses. Replication continues automatically once there is sufficient delegated space.

3. If there is not enough free replica reserve for the volume data, the primary group deletes the oldest replica to free space for the new replica. For example, if the data transfer consisted of 5 GB of new data, the replica reserve must have 5 GB of free space to store this data.

   Because each replica is a representation of volume data at the time the replication started, only the data that differentiates the oldest replica from the other replicas is deleted. Therefore, multiple replicas might be deleted to free sufficient space.

The most recent complete replica is never deleted automatically, ensuring that you always have a viable copy of volume data on the secondary group.

Note: If you cannot free enough replica reserve for the volume data by deleting replicas, the replication pauses, and the primary group generates an event message, indicating the replica reserve percentage required to complete the replication.

4. The primary group copies the volume data to replica reserve, decreasing the amount of free replica reserve. For example, if the replication transferred 5 GB of new data, free replica reserve decreases by 5 GB.

   When each replication completes, the new replica becomes the most recent complete replica.

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 you need depends on the replica reserve requirements for all the replicated volumes from a partner.

See *Guidelines for sizing replica reserve for a volume* on page 12-14 and *Guidelines for sizing delegated space* on page 12-15.

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%) 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. If there is insufficient free replica reserve for a new replica, the oldest replicas are deleted to increase free space.

Note: Replica reserve can increase automatically or by administrator action only if free delegated space is available. See *Guidelines for sizing delegated space* on page 12-15.

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* on page 12-6.

Guidelines for sizing replica reserve are as follows:

- **Recommended value**. Dell recommends that you specify 200% for the replica reserve. Specifying 200% guarantees that you can store at least two replicas, assuming there is sufficient delegated space 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%.

To obtain an appropriate replica reserve value, estimate the average volume changes that occur between replication operations. Then, use this calculation, where 105% is the minimum replica reserve value:

```
105% + [change rate x (number of replicas to keep -1)]
```

For example, if you estimate that at most 20% of the volume changes between replication operations, and you want to keep three replicas, specify 145% for the replica reserve value.

When replication is ongoing, 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 less 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.

Ideally, you should request from the secondary group administrator only enough delegated space to store the desired number of replicas for each volume.

Guidelines for sizing delegated space are as follows:

- **Recommended value**. Add together the maximum replica reserve space requirements for all the primary group volumes that 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. See *Modifying space delegated to a partner* on page 12-20.

- **Space-efficient value**. You might want to request delegated space that is less than the recommended value described above.

  Initially, replica reserve is not fully allocated. Instead, it increases automatically, based on volume usage. This enables you to over-provision delegated space. When you over-provision delegated space, the total maximum replica reserve space for all the partner volumes exceeds delegated space.

  Warning: If you over-provision delegated space, a volume's replica reserve might not be able to increase automatically or through administrative action, preventing the replication operation from completing.

For example, assume you are replicating five volumes, and the maximum combined replica reserve needed is 70 GB. If you allocate 50 GB for delegated space, delegated space is over-provisioned by 20 GB. If you specify 70 GB 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.

About replication partners

Before you can replicate volume 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.

Partners use port 3260 for replication activity.

When 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* on page 12-8.
- The secondary group must have enough free space to delegate to the primary group. See *Guidelines for sizing delegated space* on page 12-15.
- The groups must have network connectivity. The link between the 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 the same PS Series firmware version. 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.

Replication partner attributes

When you configure a replication partner, you specify values for the attributes described in Table 12-1. The first column lists attributes, and the second describes them. You can also modify the partner configuration and change the attribute settings.

Table 12-1: 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.
Description	Optional description for the partner.
Contact information	Optional contact information for the partner administrator: • Name – Up to 63 alphanumeric (ASCII) characters, including spaces. • E-mail address – Up to 31 alphanumeric (ASCII) characters, including spaces, the "@" sign, dots, dashes, and underscores. • Phone numbers – Up to 31 alphanumeric (ASCII) characters, including spaces, dots, dashes, and parentheses.
Two passwords	Passwords for mutual authentication. Each partner supplies a password to the other partner, which validates the password. Passwords are case sensitive and can consist of up to 16 alphanumeric 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* on page 12-11.

Configuring replication partners

After you obtain the replication partner attributes described in Table 12-1, you can configure the two groups as replication partners.

You must:

- Log in to the primary group (where the volume is located) and configure the secondary group as a replication partner.
- Log in to the secondary group (where the replicas are stored) and configure the primary group as a replication partner. Make sure you delegate space to the primary group.

Note: Password or configuration problems between partners do not occur 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.

On each group:

1. Click `Replication` and then `Configure partner`.
2. In the Configure Replication Partner – Identification dialog box, specify:
   - Group name. Note that group names are case sensitive.

- Group IP address.
- Optional description for the partner.

Then, click `Next`.

3. In the Configure Replication Partner – Contact dialog box, enter the optional name, e-mail address, and phone number or mobile number for the partner administrator. Then, click `Next`.
4. In the Configure Replication Partner – Authentication dialog box, enter the passwords that the partners use for mutual authentication:

   - Specify a password in the `Password for partner` field.

     When you configure the other group as a replication partner, you must specify this password in the `Password obtained from partner` field.

     For example, if you specify the password `123abc123` in the `Password for partner` field, specify `123abc123` in the `Password obtained from partner` field when configuring the other partner.

   - Specify a password in the `Password obtained from partner` field.

     When you configure the other group as a replication partner, you must specify this password in the `Password for partner` field.

     For example, if you specify the password `abc123abc` in the `Password obtained from partner` field, specify `abc123abc` in the `Password for partner` field when configuring the other partner.

   Then, click `Next`.

5. In the Configure Replication Partner – Delegated Space dialog box, optionally, select the storage pool and enter the amount of space the group delegates to the partner. Then, click `Next`.
6. In the Configure Replication Partner – Summary dialog box, review the configuration. If it is acceptable, click `Finish`.To make changes, click `Back`.

When you have configured both replication partners, you can replicate volume data.

Displaying replication partners

Click `Replication` and then `Replication Partners`.

The Replication Partners panel appears, providing the following information for each partner:

- The top entry shows whether outbound replication from the group to the partner is enabled or paused, any space that the partner delegated to the group, and free delegated space.
- The bottom entry shows whether inbound replication from the partner to the group is enabled or paused, any space that the group delegated to the partner, and free delegated space.

See the online help for information about the data fields and options.

Displaying the replication configuration for a partner

Click `Replication` and then select the partner name. The Replication Partner Summary window appears containing the following panels:

- General partner information panel – Provides information about the replication partner:
  - Group name and IP address
  - Partner information and status
- Replication status panel – Provides information about inbound replication, outbound replication, and unmanaged space:
  - Whether replication is paused or not
  - Delegated and unmanaged space
  - Volumes and collections replicated
- Replication progress panel – Provides information about in-progress replication operations:
  - Volume name and direction of transfer
  - Timestamp of transfer and transfer status
  - Amount of data transferred

See the online help for information about the data fields and options.

Modifying 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. See *Replication partner attributes* on page 12-17.

Note: Replication partner changes you make on the secondary group are not updated on the primary group until the next replication.

Modifying a partner group name or IP address

1. Click `Replication`, then select the partner, then click `Modify settings`, and then click the `General` tab.
2. In the Modify Replication Partner – General window, change the group name, IP address, or description.
3. Click `OK`.

Modifying partner contact information

1. Click `Replication`, then select the partner, then click `Modify settings`, and then click the `Contact` tab.
2. In the Modify Replication Partner - Contact window, change the contact name, e-mail address, or phone numbers.
3. Click `OK`.

Modifying 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`, then select the partner, and then click `Modify passwords`.
2. In the Modify Replication Partner - Passwords window, change each password.
3. Click `OK`.

Modifying space delegated to a partner

Restriction: You cannot decrease the space delegated to a lower capacity than is currently used to store the partner's replicas.

To modify the space delegated to a partner:

1. Click `Replication`, then select the partner, and then click `Modify delegated space`.
2. In the Modify Delegated Space dialog box, enter the new delegated space value or change the delegated space pool.

   The Pool Space table shows how pool space is currently used and how much space is used after the change. If the new delegated space exceeds the capacity of the pool, the color of the table cell showing free pool space changes to red. Changing the pool moves all the replicas and any recovery volumes for the partner to the new pool.
3. Click `OK`.

Deleting 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. Then, the delegated space becomes free pool space. However, 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, do one of the following before you delete the partner:

- 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. Click `Replication`, then select the partner, and then click `Pause inbound`.
2. Click `Delete partner.`
3. Confirm that you want to delete the partner.

Displaying inbound and outbound replication

Click `Replication` and then expand the partner name. From the Replication Partner Status window, you can:

- Click `Inbound Replicas` to display partner replicas stored in the group. Select an individual replica to display detailed information.
- Click `Inbound Replica Collections` to display partner replica collections stored in the group. Select an individual replica collection or replica to display detailed information.
- Click `Outbound Replicas` to display replicas of group volumes stored on the partner. Select an individual replica to display detailed information.
- Click `Outbound Replica Collections` to display replicas of group volume collections stored on the partner. Select an individual replica or replica collection to display detailed information.

Note: Individual replicas are identified by the date and time the replication operation started. The time stamp for the primary group is based on its time zone; likewise, the time stamp for the secondary group is based on its time zone.

The Inbound Replicas and Inbound Replica Collections windows show the following:

- The Delegated Space panel shows the usage of the space that the group delegated to the partner.
- The Local Replicas panel shows each replicated volume or volume collection, the volume replicas, the replica reserve size, the replication status, and the replica status.

The Outbound Replicas and Outbound Replica Collections windows show the following:

- The Replica Space panel shows the usage of the space the partner delegated to the group.
- The Remote Replicas panel shows each replicated volume or volume collection, the volume replicas, the replica reserve size, the replication status, and the replica status.

In the Remote Replicas panel, click `Replication history` to display details about the last 10 replication operations for each volume, including the duration of the replication and the size and speed of the data transfer. The duration includes the amount of time during which replication was paused or the network was down.

Displaying inbound replica collections

To display all inbound replica collections, click `Replication`, then expand the partner name, and then expand `Inbound Replica Collections`.

To display an individual replica collection, click `Replication`, then expand the partner name, then expand `Inbound Replica Collections`, and then select the volume collection name.

The Inbound Replicas Collections window appears, containing the following panels:

- Delegated space panel – Provides information about the space the group delegated to the selected partner:
  - Space delegated, used, and free
  - Failback replica space
  - Inbound replica status
- Inbound replicas panel – Provides information about inbound replicas and replica status.

See the online help for information about the data fields and options.

Displaying all inbound replicas

Click `Replication`, then expand the partner name, and then expand `Inbound Replicas`.

The Inbound Replica Summary window appears, containing the following panels:

- Delegated Space – Provides information about the space that the group delegated to the partner.
- Inbound Replicas panel – Provides information about each replicated volume, the volume replicas, the replica reserve size, the replication status, and the replica status.

See the online help for information about the data fields and options.

Displaying individual inbound replica sets

Click `Replication`, then expand the partner name, then expand `Inbound Replicas`, and then select the replica set name. The Replica Set Status window appears, containing the following panels:

- General replica set information panel – Provides information about the inbound volume replication configuration:
  - Status and latest transfer timestamp
  - Whether a failback snapshot or thin provisioning are enabled.
  - User description and identifiers: Volume name, storage pool, and replication partner
  - Requested and free reserve
- Replicas panel – Provides information about the replicas created for the selected volume.

Note: Select the column headers in the GUI to sort the table. By default, the table is sorted by the Created date.

See the online help for information about the data fields and options.

Displaying an inbound template replica set

To display the inbound replica set for a template volume, click `Replication`, then expand the partner name, then expand `Inbound Replicas`, and then select the replica set name. The Template Replica Set Status window appears. For a template volume replica set, this window contains two tabs.

Template replicas tab

In the Replica Set Status window, the Template Replicas tab provides information about the inbound replica set for the template volume and includes the following panels:

- Template replicas panel – Provides information about template volume replication configuration.
  - Size, status and latest transfer timestamp
  - Whether a failback snapshot or thin provisioning are enabled
  - Storage pool, and replication partner
  - Requested and free reserve
- Replicas panel – Provides information about the replicas created for the selected volume.

Thin clone replica sets tab

The Thin Clone Replica Sets tab shows:

- Number of thin clone replica sets attached to the template replica set.
- Number of promoted thin clone replica sets (that is, recovery thin clones) attached to the template replica set.
- Thin Clone Replica Sets panel – Provides information about each thin clone replica set or promoted thin clone replica set, including:
  - Name, reported size and status
  - Replication partner and status
  - iSCSI Connections

See the online help for information about the data fields and options.

Displaying all outbound replica collections

Click `Replication`, then expand the partner name, and then click `Outbound Replica Collections`. The Outbound Replica Collections window appears, containing the following panels:

- Remote delegated space panel – Provides information about capacity usage for delegated space on the selected replication partner.
- Remote replicas panel – Provides information about the replicas stored on a partner.

See the online help for information about the data fields and options.

Displaying individual outbound replica collections

Click `Replication`, then expand the partner name, then expand `Outbound Replica Collections`, and then select the volume collection. The Replication of Collection *name* panel appears, containing the following panels

Note: You can also display information about volume collection replication by clicking `Volumes` in the lower-left corner, then expanding `Volume Collections`, then selecting the collection name, and then clicking the `Replicas` tab.

- Replication summary panel – Provides information about the volume collection replication configuration, including replication partner, and schedule status.
- Remote replicas panel – Provides information about the replicas stored on a replication partner.

See the online help for information about the data fields and options.

Displaying all outbound replicas

Click `Replication`, then expand the partner name, and then click `Outbound Replicas`.

The Outbound Replicas Summary window appears, containing the following panels:

- The remote delegated space panel – Provides information about the space the selected partner delegated to the group.
- Remote replicas panel – Provides information about replicas stored on the partner.

In the Remote Replicas panel, select `Replication history` to display details about the last 10 replication operations for each volume, including the duration of the replication and the size and speed of the data transfer. The duration includes the amount of time during which replication was paused or the network was down.

See the online help for information about the data fields and options.

Displaying the outbound replication of an individual volume

Click `Replication`, then expand the partner name, then expand `Outbound Replicas`, and then select the volume. The Replication of volume *name* window appears, containing the following panels:

Note: You can also display information about volume replication by clicking `Volumes` in the lower-left corner, then expanding `Volumes`, then selecting the volume name, and then clicking the `Replication` tab.

- Replication summary panel – Provides information about the replication configuration for the volume:
  - Replication partner
  - Replica and local reserve
  - Failback snapshot an baseline
  - Schedule and status
- Remote replicas panel – Provides information about each replica:
  - Click `Replication history` to display details about the last 10 replication operations for the volume.
  - Select `Thin clone replicas` (only applicable if the outbound replica is for a template volume) to display information about any replicated thin clones attached to the template volume.

Volume replication configuration attributes

Table 12-2 describes the attributes you set when configuring the volume for replication. The first column lists their attributes and the second describe them. You can modify the replication configuration and change the attribute values.

Table 12-2: Volume Replication Configuration Attributes

Attribute	Description
Replication partner	Partner that stores the volume replicas. The partner must have space delegated to the group. See *Configuring replication partners* on page 12-17.
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. See *About local replication reserve* on page 12-8.
Borrow space setting	Enables you to borrow from free pool space if the local replication reserve size is inadequate. **Note:** To enable the borrow space setting or to borrow space, the pool must have at least 10% free pool space.
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 the volume replicas. See *About delegated space and replica reserve* on page 12-11.

Configuring 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.

Restriction: You cannot configure a thin clone for replication until you replicate the template volume to which the thin clone is attached.

To configure a volume for replication:

1. Gather the volume replication attributes. See *Volume replication configuration attributes.*
2. Click `Volumes` in the lower-left panel, then expand `Volumes`, then select the volume, and then click `Configure replication.`
3. In the Configure Volume Replication – General Settings window:
   - Select the replication partner.
   - Specify the replica reserve percentage.
   - Specify the local replication reserve percentage and whether to allow borrowed pool space.

   Then, click `Next`.
4. In the Configure Volume Replication – Advanced Settings window, select whether to keep the failback snapshot and click `Next`.

Because 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.

5. In the Configure Volume Replication – Summary window, review the information and click `Finish` if the configuration is correct. Click `Back` to make changes.

When you complete the volume replication configuration, you can choose to create a replica. You can also choose to perform the replication by using manual transfer replication. See *About replicas* on page 12-1.

Modifying volume replication configuration settings

You can modify the settings described in *Volume replication configuration attributes*. Changes are not applied until the next replication.

To modify the settings, click `Volumes`, then expand `Volumes`, then select the volume, then click `Modify replication settings`, and then click the `General` tab.

In the Modify Volume Replication Settings – General dialog box, you can change:

- Replication partner. The replicas remain on the original partner. The next replication to the new partner transfers the full volume contents.
- Replica reserve percentage.

  Note: The space currently used to store replicas represents the lower limit for replica reserve. You cannot decrease replica reserve below this limit.

- Local replication reserve percentage.
- Borrow space setting.

To change the failback snapshot setting, click the `Advanced` tab. In the Modify Volume Replication Settings – Advanced dialog box, select or deselect `Keep failback snapshot`.

Note: If you select `Keep failback snapshot`, you must create a replica to establish the failback snapshot.

Configuring 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. See *Configuring a volume for replication* on page 12-25.

1. Click `Volumes`, then expand `Volume Collections`, then select the collection, and then click `Configure replication`.
2. In the Volume Collection – Modify Replication Settings dialog box:
   - Select the replication partner for the volume collection.
   - Make sure each volume is configured to replicate to the partner selected above.

     If a volume is not configured for replication, click the `not replicated` link and configure the volume.

See *Configuring a volume for replication* on page 12-25.
If a volume is configured to replicate to a different partner, click the partner name link to modify the volume replication configuration and change the partner. See *Modifying volume replication configuration settings* on page 12-26.

3. Click `OK`.

Modifying volume collection replication configuration settings

You can modify the replication configuration of a volume collection or the volumes in the collection.

1. Click `Volumes`, then expand `Volume Collections`, then select the collection, and then click `Modify replication settings.`
2. In the Volume Collection – Modify Replication Settings dialog box, you can change the partner name for the collection. You can also click the partner name next to a volume to change the replication configuration for the volume. See *Modifying volume replication configuration settings* on page 12-26.
3. Click `OK`.

Disabling 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.

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.

Note: When you disable replication on a volume, the delegated space on the secondary group that is storing the replicas becomes unmanaged space; that is, the space cannot be managed from the primary group. If you do not need the replicas, log into the secondary group and delete the replica set.

Restriction: 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`, then expand `Volumes`, then select the volume name, and then click `Disable volume replication.`
2. Confirm that you want to disable replication on the volume.

To disable replication for a volume collection:

1. Click `Volumes`, then expand `Volume Collections`, then select the collection name, and then click `Disable volume replication.`
2. Confirm that you want to disable replication on the volume collection.

Creating a replica

The first time 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.

Note: 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. See *About manual transfer replication* on page 12-3.

Prerequisite: Set up the replication partners and configure the volume or volume collection for replication.

Prerequisite: You must replicate a template volume before replicating any of its thin clones.

Restriction: You can replicate a template volume only one time.

See *Configuring replication partners* on page 12-17, *Configuring a volume for replication* on page 12-25.

1. Click `Volumes`, then expand `Volumes`, then select the volume name, and then click `Create replica now.`
2. In the Create Replica dialog box:
   - If you want to use manual transfer replication to perform the replication, select `Perform manual replication`.
   - Click `Yes` to start the replication.
3. Monitor replication to make sure replicas complete in a timely manner. See *Displaying replication activity and replicas for a volume* on page 12-28.

If replication operations take longer than expected, make sure you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link can cause long replication times.

Displaying replication activity and replicas for a volume

1. Click `Volumes` in the lower-left panel, then expand `Volumes`, then select the volume name, and then click the `Replication` tab. In the Volume Replication window, the Replication Summary panel shows:
   - Current size of the replica reserve.
   - Local replication reserve.
   - Whether you are keeping the failback snapshot and, if so, the failback baseline timestamp.
   - The amount of data that must be transferred for the next or current replication.
   - Any replication schedules and the next scheduled replication.

   The Remote Replicas panel shows:
   - Replica set and individual replicas for the volume. Each replica is identified by the date and time the replication operation started.
   - Number of complete replicas and the current size of the replica reserve.

- Replication operation status and replica status.

2. In the Remote Replicas panel, click `Replication History` to show details about each replication operation:
   - Time the operation started.
   - Replication partner.
   - Total replication time. The duration time includes the amount of time during which replication was paused or the network was down.
   - Amount of data transferred.
   - Data transfer speed.
   - Status of the replication.

Replicating 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.

Requirement: Before you can replicate a volume collection, the volumes in the collection and the volume collection must be configured to replicate to the same partner. See *Configuring a volume for replication* on page 12-25 and *Configuring a volume collection for replication* on page 12-26.

1. Click `Volumes`, then expand `Volume Collections`, then select the collection, and then click `Create replica now.`
2. In the Create Replica Collection dialog box:
   - If you want to perform the replication using manual transfer replication, select `Perform manual replication`. See *About manual transfer replication* on page 12-3.
   - Click `Yes` to start the replication.

You should monitor replication to make sure replicas complete in a timely manner. See *Displaying replication activity and replicas for a volume collection* on page 12-29.

If replication operations are taking longer than expected, make sure you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link can cause long replication times.

Displaying replication activity and replicas for a volume collection

Click `Volumes`, then expand `Volume Collections`, then select the collection name, and then click the `Replicas` tab.

In the Volume Collection Replicas window, 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.

Using schedules to create replicas

Schedules enable you to create replicas of a volume or all the volumes in a collection on a regular basis.

Note: Schedules apply only to network replications. Scheduled replications do not run until any in-process manual transfer replications are complete.

Restriction: You cannot schedule replicas for a volume template.

Before you set up a replication schedule:

- Configure a partner. See *Configuring replication partners* on page 12-17.
- Make sure the volume or volume collection is configured for replication. See *Configuring a volume for replication* on page 12-25 and *Configuring a volume collection for replication* on page 12-26.

To set up a replication schedule, see *Scheduling volume operations*.

Pausing and resuming 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, click `Volumes`, then expand `Volumes`, then select the volume name, and then click `Pause volume replication.`

To resume replication for a volume, click `Volumes`, then expand `Volumes`, then select the volume name, and then click `Resume volume replication.`

Pausing and resuming replication to or from a partner

You can pause and then resume replication as needed. Some operations require temporarily pausing replication.

Pausing and resuming outbound replication

- To pause outbound replication to a partner:

  Click `Replication`, then expand the partner name, and then click `Pause outbound`.
- To resume outbound replication to a partner:

  Click `Replication`, then expand the partner name, and then click `Resume outbound`.

Pausing and resuming inbound replication

- To pause inbound replication from a partner:

  Click `Replication`, then expand the partner name, and then click `Pause inbound`.

- To resume inbound replication from a partner:

  Click `Replication`, then expand the partner name, and then click `Resume inbound`.

Cancelling a volume replication

You can cancel an in-progress volume replication.

Note: To temporarily stop volume replication instead of cancelling it, pause the replication. See *Pausing and resuming replication of a volume* on page 12-30.

Click `Volumes`, expand Volumes, then select the volume name, and then click `Cancel replication`.

Cloning an inbound replica

Cloning enables data access in a replica, with no impact on the replication configuration or the replica.

You can clone an inbound replica to create a new volume on the secondary group. The new volume has the same reported size, is the same type (standard volume, template volume, or thin clone volume), and has the same data as the original volume at the time you created the replica.

The new volume is located in the same pool as the replica. After the clone operation completes, the replica still exists, and replication continues as usual.

Note: 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.

1. Click `Replication`, then expand the replication partner, then expand `Inbound replicas`, then select the replica set, and then click `Clone replica`.
2. In the Clone Replica dialog box, select the timestamp of the replica.
3. In the Clone Volume Replica – Volume Settings dialog box, specify the new volume name and (optional) description. Then, click `Next`.
4. In the Clone Volume Replica – Space dialog box, change the thin provisioning settings and the snapshot reserve value, and click `Next`.
5. In the Clone Volume Replica – iSCSI Access dialog box, specify:

   - Access controls. See *About iSCSI target access controls*.
   - Permission, either read-only or read-write.
   - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume and its snapshots. See *Multi-host access to targets*.

   Then, click `Next`.
6. In the Clone Volume Replica – Summary dialog box, review the information and click `Finish` if satisfactory. Click `Back` to make changes.

The new volume appears in the list of volumes in the far-left panel.

Deleting outbound replica sets or replicas

Deleting a replica set disables replication on the volume. If you re-enable replication on the volume, the first replication is a complete transfer of volume data.

1. Click `Volumes`, then expand `Volumes`, then select the volume name, and then click the `Replication` tab.
2. In the Volume Replication window, click `Volume replicas` in the Remote Replicas panel.
3. To delete the replica set, select the replica set and click `Delete replica set`.

   To delete a replica, select the replica, and click `Delete replica`.
4. Confirm that you want to delete the replica set or replica.

Deleting 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.

If you delete a replica from a replica collection, the replica collection 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 re-enable replication on the volume collection, the first replication of each volume is a complete transfer of volume data.

1. Click `Volumes`, then expand `Volume Collections`, then select the collection name, and then click the `Replicas` tab.
2. In the Volume Collection Replicas window, to delete the replica collection set, select it in the Remote Replicas panel and click `Delete replica collection set`.

   To delete a replica collection, expand the replica collection set, select the replica collection, and click `Delete replica collection`.

   To delete a replica from a replica collection, expand the replica collection, select the replica, and click `Delete replica`.
3. Confirm that you want to perform the delete operation.

Deleting inbound replica sets or replicas

If the primary group is not available, you can delete replicas and replica sets when 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. Deleting a replica set disables replication on the volume.

Recommendation: Dell recommends that you delete replicas when logged in to the primary group. See *Deleting outbound replica sets or replicas* on page 12-32.

To delete an inbound replica or replica set:

1. Click `Replication`, then expand the partner name, and then click `Inbound Replicas.`
2. If you want to delete an inbound replica set, click `Pause inbound`.
3. In the Inbound Replicas panel:

   To delete a replica set, select the replica set and click `Delete replica set`.

   To delete a replica, expand the replica set, select the replica, and click `Delete replica`.
4. Confirm that you want to delete the replica or replica set.
5. If you paused inbound replication, click `Resume inbound`.

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.

Deleting inbound replica collection sets, replica collections, or replicas

You can delete an inbound 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 and disables replication on the volume collection.

Recommendation: Dell recommends that you delete replicas when logged in to the primary group. See *Deleting outbound replica collection sets, replica collections, or replicas* on page 12-32.

1. Click `Replication`, then expand the partner name, and then click `Inbound Replica Collection`.
2. If you are deleting an inbound replica collection set, click `Pause inbound`.
3. To delete a replica collection set, select it in the Inbound Replicas panel and click `Delete replica collection set`.

   To delete a replica collection, expand the replica collection set, select the replica collection, and click `Delete replica collection`.

   To delete a replica from a replica collection, expand the replica collection, select the replica, and click `Delete replica`.
4. Confirm that you want to perform the delete operation.
5. If you paused inbound replication, click `Resume inbound`.

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.

13 Data recovery

If you replicate a volume to a partner (see Chapter 12, *Volume replication*), you can recover volume data on the partner. In addition, you might be able to fail over to the partner and later fail back to the original group.

About data recovery

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. See Chapter 12, *Volume replication*.

If the volume becomes unavailable—either temporarily or permanently—you can recover the data from the partner.

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.

The method for recovering data depends on the state of the groups and your specific data recovery requirements. See *Data recovery procedures*.

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. See *Cloning an inbound replica* on page 12-31.

In most situations in which you must 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.

See *Failing over and failing back a volume* on page 13-2 for more details.

Data recovery procedures

Table 13-1 describes common data recovery procedures.

Table 13-1: Data Recovery Procedures

Goal	Procedure	Considerations	Reference
Fail over and fail back a volume. Use this method if the volume is unavailable due to a failure or maintenance.	1. Promote the replica set to a recovery volume. 2. Demote the volume to a failback replica set. 3. Replicate the recovery volume. 4. Demote the recovery volume to a replica set. 5. Promote the failback replica set to a volume.	If the failback snapshot is not available on the primary group, you must replicate the full recovery volume, instead of just the changes.	*Failing over and failing back a volume* on page 13-2
Make a temporary copy of volume data available on the primary group.	1. Promote the replica set to a recovery volume. 2. Perform the operation on the recovery volume. 3. Demote the recovery volume to an inbound replica set.	This method assumes you do not want to preserve writes made to the recovery volume.	*Making a temporary volume available on the secondary group* on page 13-12
Permanently switch partner roles in a volume replication configuration.	1. Permanently promote the replica set to a volume. 2. Permanently demote the volume to an inbound replica set. 3. Configure the volume to replicate from the new primary group to the new secondary group.	If the failback snapshot does not exist, the first replication after the role switch is a complete copy of the volume data.	*Permanently switching partner roles* on page 13-12
Permanently host the volume on the group that was the secondary group.	1. If the volume is available, set the volume offline. 2. Permanently promote the replica set to a volume. 3. Optionally, replicate the new volume.	The first replication of the new volume is a complete copy of the volume data.	*Permanently promoting a replica set to a volume* on page 13-14

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.

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. See *Promoting an inbound replica set to a 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:

   - Demote the original volume to a failback replica set on the primary group.
   - 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.

See *Replicating a recovery volume to the primary group*.

3. When you are ready to fail back to the primary group, use the Failback to Partner operation to:
   - Set the recovery volume offline.
   - Perform a final replication to synchronize the volume data across both groups.
   - Demote the recovery volume to an inbound replica set.
   - 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.

   See *Failing back to the primary group*.

At this point, initiators can connect to the volume on the primary group and replication can continue as usual. By default, the failback baseline is reestablished.

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 13-1 to Figure 13-5. Figure 13-1 shows the replication configuration, where GroupA is replicating Volume1 to GroupB.

Figure 13-1: No Failure – Data Is Available

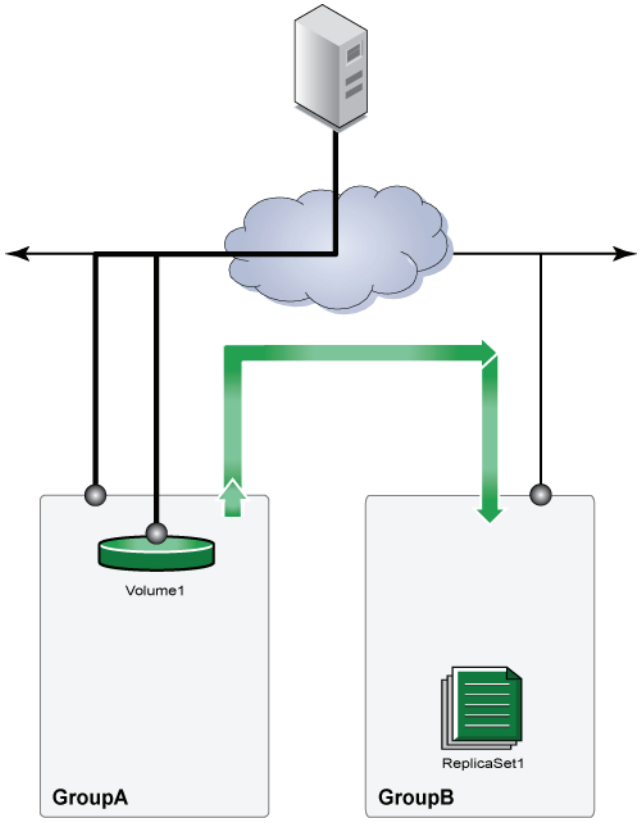

Figure 13-2 shows the replication configuration after a failure in the primary group (GroupA).

Figure 13-2: Primary Group Failure – Data Is Not Available

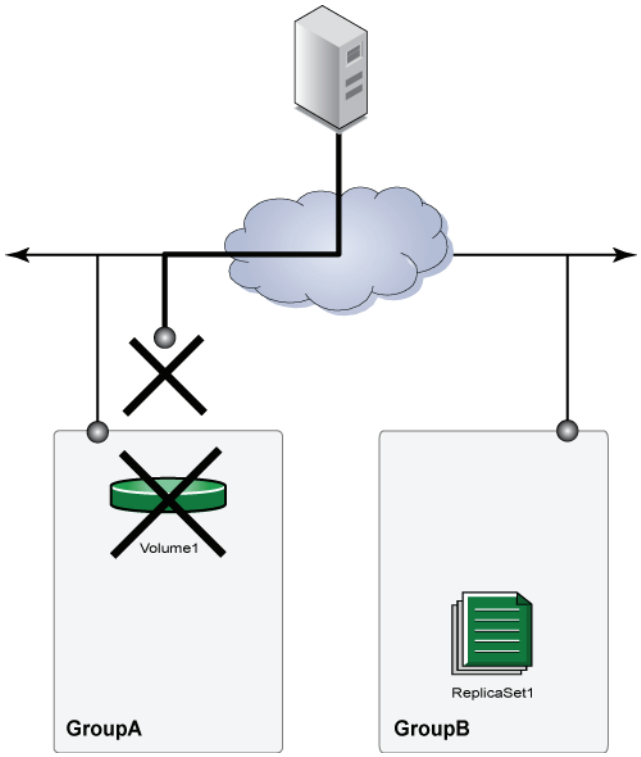

Figure 13-3 shows the first step in recovering data on the secondary group, which is to fail over the volume to the secondary group. To do this, promote the inbound replica set to a recovery volume and snapshots. The recovery volume contains the volume data represented by the most recent complete replica. Users can connect to the recovery volume to resume access to volume data.

Figure 13-3: Step 1: Failover to the Secondary Group – Data Is Available

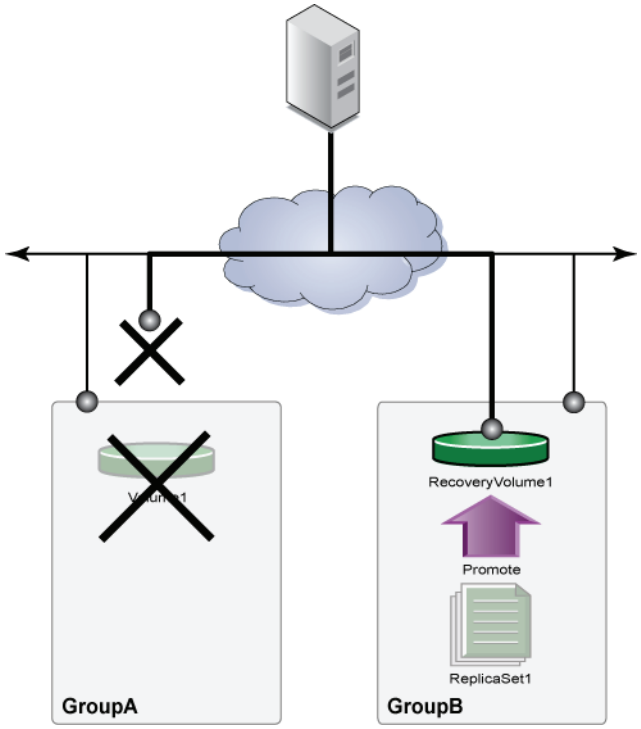

Figure 13-4 shows the second step in recovering data—replicate to the primary group. When the primary group is available:

- 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 13-4: Step 2: Replicate to the Primary Group – Data Is Available and Protected

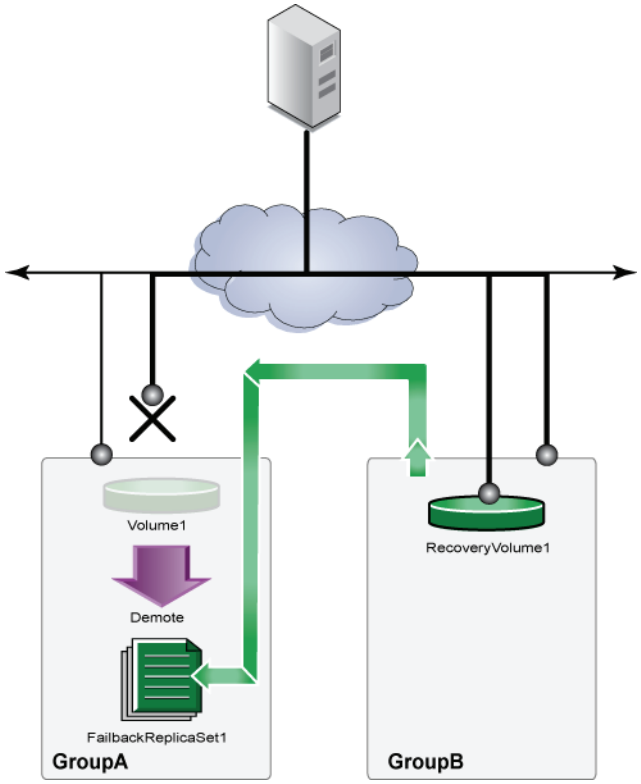

Figure 13-5 shows the final step in recovering data—fail back to the primary group. To fail back to the primary group:

- 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 Replicate 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 13-5: Step 3: Fail Back to the Primary Group

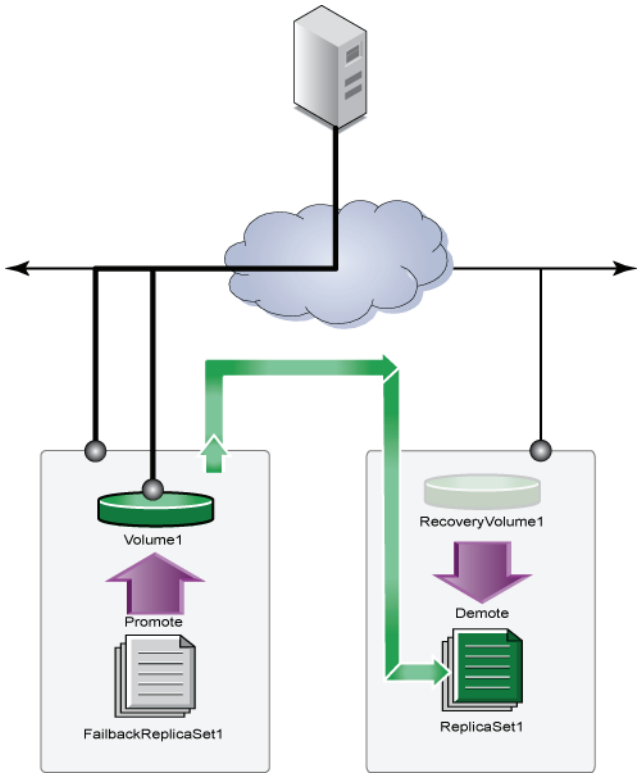

Promoting an inbound replica set to a recovery volume

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 access 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 chose 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.

With some exceptions, all volume operations apply to a recovery volume. See *Recovery volume restrictions*.

Restriction: You cannot convert a recovery template to a standard volume. You must first You must first make the promotion permanent. You cannot detach a recovery thin clone.

1. Click `Replication`, then expand the replication partner, then expand `Inbound Replicas`, then select the replica set, and then click `Promote to volume`.
2. Confirm that you want to pause inbound replication from the partner. Replication resumes automatically for all *other* volumes after the replica set is promoted.
3. In the Promote Replica Set – Volume Options dialog box, specify the following:

   - Whether you want to set the recovery volume online or offline. Set the volume online if you want initiators to connect to it.
   - Whether you want to retain the iSCSI target name of the original volume.
   - Whether you want to keep the ability to demote to replica set. Unless you are permanently promoting the replica set, make sure you keep this ability.

   Then, click `Next`.
4. In the Promote Replica Set – iSCSI Access, specify the following:

   - Conditions that a computer must match to connect to the recovery volume. Specify a CHAP user name, IP address, or iSCSI initiator name. See *About iSCSI target access controls* on page 8-1.
   - Recovery volume permission, either read-only or read-write.
   - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the recovery volume. See *Multi-host access to targets* on page 8-6.

   Then, click `Next`.
5. In the Promote Replica Set – Summary dialog box, review the information and click `Finish` if satisfactory. 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.

Where to go next

- Connect to the recovery volume. See *Connecting initiators to iSCSI targets* on page 8-7.
- Go to the next step in the failover and failback process when the original volume on the primary group becomes available. See *Replicating a recovery volume to the primary group* on page 13-9.
- To reverse the inbound replica set promotion and cancel the failover, demote the recovery volume to an inbound replica set:

  Click `Volumes`, then expand `Volumes`, then select the recovery volume name, and then click `Demote to replica set`.

- If the original volume becomes permanently unavailable, you can make the inbound replica set promotion permanent. See *Making an inbound replica set promotion permanent* on page 13-13.

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 access 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)

In addition, you cannot delete a recovery template volume if there are still recovery thin clone volumes, thin clone replica sets, or permanently promoted thin clone replica sets attached to the volume.

See Chapter 9, *Basic volume operations* for information about modifying 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. This is 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 4.0 or greater. If you do not meet this requirement, you must perform the steps individually, as described in *Manually performing the replicate to partner operation* on page 13-16.

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.

1. Obtain the name and password for a group administrator account on the primary group.
2. Click `Volumes`, expand `Volumes`, select the recovery volume, and then click `Replicate to partner`.

3. In the Replicate Recovery Volume dialog box:
   - Specify the group administrator account name and password.
   - Select whether to perform the replication by using manual transfer replication. See *About manual transfer replication* on page 12-3.
   - Select whether to save the primary group administrator account name and password for future use in the current GUI session.
4. Click `OK`.

To monitor the Replicate to Partner operation and make sure all tasks complete, open the Alarms panel at the bottom of the GUI window and click the `Failback Operations` tab. If an individual task fails, you must correct the problem and then retry the task. See *Handling a failed operation* on page 13-15.

Note: If you chose 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.

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 next

- To create more replicas, select the recovery volume and click `Create replica now`. You can also configure replication schedules on the recovery volume.
- When you are ready to fail back, see *Failing back to the primary group*.
- 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. On the primary group, click `Replication`, expand the replication partner, expand `Inbound Replicas`, select the failback replica set, and click `Change storage pool`.
  2. Select the new pool and click `OK`.
- If you do not want to return to the original replication configuration or switch roles, you can make the inbound replica set promotion permanent and then delete the replica set:

  1. On the primary group, click `Replication`, then expand the replication partner, then expand `Inbound Replicas`, then select the failback replica set, and then click `Convert to replica set`.

     The replica set continues to be shown in the Replication Partner – Inbound window, but it is no longer a failback replica set.
  2. Click `Pause inbound`.
  3. Select the replica set and then click `Delete replica set`.

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. If you later promote the failback replica set to a volume, the volume belongs to the new pool.

1. On the primary group, click `Replication`, expand the replication partner, expand `Inbound Replicas`, select the failback replica set, and click `Change storage pool`.
2. Select the new pool and click `OK`.

Failing back to the primary group

When you want 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. See *Replicating a recovery volume to the primary group*.

Restriction: You cannot fail back a template volume.

Restriction: The Failback to Primary operation is available only if the primary group and the secondary group are running PS Series Firmware Version 4.0 or greater. If you do not meet this requirement, you must perform the steps individually, as described in *Manually performing the failback to primary operation* on page 13-16.

1. Obtain the name and password for a group administrator account on the primary group.
2. Click `Volumes`, then expand `Volumes`, then select the recovery volume, and then click `Failback to primary`.
3. Confirm that you want to set the recovery volume offline.
4. In the Failback Recovery Volume dialog box:
   - Specify the group administrator account name and password.
   - Select whether to perform the replication by using manual transfer replication. See *About manual transfer replication* on page 12-3.
   - Select whether to save the primary group administrator account name and password for future use in the current GUI session.

   Then, click `OK`.

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.

To monitor the Failback to Primary operation and make sure all tasks complete, open the Alarms panel at the bottom of the GUI window and click the `Failback Operations` tab. If an individual task fails, you must correct the problem and then retry the task. See *Handling a failed operation* on page 13-15.

Note: If you chose 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 Failback to Primary operation continues automatically.

When the Failback to Primary operation completes, on the secondary group, the recovery volume disappears from the list of volumes, and the inbound replica set reappears in the list of inbound replica sets. On the primary group, the failback replica set disappears from the list of inbound replica sets, and the volume reappears in the list of volumes.

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. This 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 *Failing over and failing back a volume*.

1. Promote the replica set to a recovery volume. See *Promoting an inbound replica set to a recovery volume*.

   Make sure you select the option that enables you to demote the recovery volume.

2. Perform the desired operation on the recovery volume. See *Recovery volume restrictions* on page 13-9.

3. Demote the recovery volume to an inbound replica set:

   Click `Volumes`, then expand `Volumes`, then select the recovery volume, and then click `Demote to replica set`.

At this point, you can resume replicating the volume.

Permanently 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.

1. On the primary group:

   a. Make sure the volume replication configuration includes keeping the failback snapshot. See *Modifying volume replication configuration settings* on page 12-26.

   b. Set the volume offline. See *Setting a volume offline or online* on page 9-13.

   c. Perform a final replication. This synchronizes volume data across the primary group and the secondary group. See *Creating a replica* on page 12-28.

2. On the secondary group:

   a. Promote the replica set to a recovery volume. Make sure you keep the ability to demote the recovery volume, in case you decide to cancel the role switch. See *Promoting an inbound replica set to a recovery volume*.

      Users can now access volume data by connecting to the recovery volume. See *Connecting initiators to iSCSI targets* on page 8-7.

   b. Replicate the recovery volume to the primary group. See *Replicating a recovery volume to the primary group*.

   c. Make the inbound replica set promotion permanent. See *Making an inbound replica set promotion permanent*.

3. On the primary group, convert the failback replica set to an inbound replica set. See *Converting a failback replica set to an inbound replica set*.

The partner role switch is complete.

Making 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.

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.

1. Click `Volumes`, then expand `Volumes`, then select the recovery volume, then click `Make promote permanent`.

2. In the Convert Recovery Volume – Volume Settings dialog box:

   - Enter a new volume name, up to 63 alphanumeric characters (including periods, dashes, and colons). A volume name must be unique in a group.
   - Enter an optional description.
   - Select the storage pool.

   Then, click `Next`.

3. In the Convert Recovery Volume – iSCSI Access dialog box, specify:

   - Access control credentials for the recovery volume. Specify a CHAP user name, IP address, or iSCSI initiator name. See *About iSCSI target access controls* on page 8-1.
   - Permission, either read-only or read-write.
   - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume. See *Multi-host access to targets* on page 8-6.

Then, click `Next`.

4. Review the information in the Convert Recovery Volume – Summary dialog box and click `Finish` if satisfactory. Click `Back` to make changes.

When the operation completes, the recovery volume is converted to a volume.

Where to go next

- If you are permanently switching partner roles, see *Converting a failback replica set to an inbound replica set* for the next step in the procedure.

Converting 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.

Note: After you convert a failback replica set to an inbound replica set, you cannot promote the inbound replica set to the original volume.

1. On the primary group, click `Replication`, then expand the partner, then expand `Inbound Replicas`, then select the failback replica set, then click `Convert to replica set`.

2. Confirm that you want to convert the replica set.

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.

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.

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.

Restriction: In some cases, you cannot permanently promote a replica set in a single operation. If you do not get the option to deselect the `Keep ability to demote to replica set` option, you must temporarily promote the replica set and then make the promotion permanent. See *Promoting an inbound replica set to a recovery volume* on page 13-7 and *Making an inbound replica set promotion permanent* on page 13-13.

Restriction: Before you can permanently promote a template replica set, you must permanently promote all the attached thin clone replica sets.

1. On the secondary group, click `Replication`, then expand the partner, then expand `Inbound Replicas`, then select the replica set name, and then click `Promote to volume`.

2. Confirm that you want to pause inbound replication from the partner.

3. In the Promote Replica Set – Volume Options dialog box:

   - Choose whether to set the volume online or offline.
   - Choose whether to retain the iSCSI target name of the original volume. This can facilitate initiator access to the volume.
   - Deselect the `Keep ability to demote to replica set` option.

   Then, click `Next`.

4. In the Promote Replica Set – Volume Settings dialog box:

   - Enter a new volume name, up to 63 alphanumeric characters (including periods, hyphens, and colons). A volume name should be unique in a group.
   - Enter an optional description.
   - Select the storage pool.

   Then, click `Next`.

5. In the Promote Replica Set – iSCSI Access dialog box:

   - Access controls for the recovery volume. Specify a CHAP user name, IP address, or iSCSI initiator name. See *About iSCSI target access controls* on page 8-1.
   - Permission, either read-only or read-write.
   - Whether to allow initiators with different iSCSI qualified names (IQNs) access to the volume. See *Multi-host access to targets* on page 8-6.

   Then, click `Next`.

6. If the information in the Promote Replica Set – Summary dialog box is satisfactory, click `Finish`. 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.

Handling a failed operation

The Replicate to Partner operation and the Failback to Primary operation consolidate multiple tasks, as documented in *Manually performing the replicate to partner operation* on page 13-16 and *Manually performing the failback to primary 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 select `Retry task`. The operation continues automatically.

Manually performing the replicate to partner operation

The Replicate to Partner operation consolidates multiple tasks. You can perform each task in the operation individually.

Requirement: You must promote the inbound replica set to a recovery volume, before performing the individual Replicate to Partner tasks. See *Promoting an inbound replica set to a recovery volume*.

1. On the primary group:
   a. Set the original volume offline. See *Setting a volume offline or online* on page 9-13.
   b. Cancel any in-progress replication. See *Cancelling a volume replication* on page 12-31.
   c. Set any snapshots for the volume offline. See *Setting a snapshot online or offline* on page 11-11.
   d. Demote the volume to a failback replica set:

      Click `Volumes`, then expand `Volumes`, then select the volume, and then click `Demote to replica set`.
2. On the secondary group:
   a. Configure the recovery volume to replicate to the primary group. See *Configuring a volume for replication* on page 12-25.
   b. Create a replica. See *Creating a replica* on page 12-28.

      You can use manual replication if a large amount of data must be transferred. See the Manual Transfer Utility *Installation and User Guide*.

Manually performing the failback to primary operation

The Failback to Primary operation consolidates multiple tasks. You can perform each task in the operation individually.

1. Perform the Replicate to Partner operation. See *Replicating a recovery volume to the primary group* on page 13-9 or *Manually performing the replicate to partner operation*.
2. On the secondary group:
   a. Disable any replication or snapshot schedules for the recovery volume.
   b. Set the recovery volume offline. See *Setting a volume offline or online* on page 9-13.
   c. Create a final replica. See *Creating a replica* on page 12-28.

      You can use manual replication if a large amount of data must be transferred. See the Manual Transfer Utility *Installation and User Guide*.
   d. Demote the recovery volume to the original inbound replica set:

      Click `Volumes`, then expand `Volumes`, then select the volume, and then click `Demote to replica set`.

3. On the primary group, promote the failback replica set to the original volume:

   Click `Replication`, expand the partner name, expand `Inbound Replicas`, select the failback replica set, and click `Promote to volume`.

Part III: Troubleshooting

14 Group event logging

A PS Series group generates events when normal group operations occur and also when significant events occur. Events enable you to track operations and also detect and solve problems before they affect performance or data availability.

In addition, you can use SNMP traps to track significant group events.

About event messages

When an event occurs in the group (for example, you create a volume or a power supply fails), the group generates an event message.

Event messages help you monitor normal operations and also identify problems before they disrupt service.

Events appear in the Monitoring window, the CLI console, and in the output of the `show recentevents` command.

Note: You can disable the display of informational messages in the GUI and CLI; however, the group still logs these events. See *Enabling or disabling the display of INFO event messages* on page 14-5.

Examples of event messages seen on the console are:

```
484:2:gigan34mem1:netmgtd:15-Mar-2010 11:25:04.310003:rcc_util.c:714:INFO:25.2.9:CLI: Login
to account grpadmin succeeded.

10:5:gigan34mem1:SP:13-Mar-2010 22:30:19.250006:emm.c:1922:WARNING:28.3.51:Warning health
conditions currently exist. Correct these conditions before they affect array operation.

Control modules are initializing. Control module failover cannot occur until the
initialization completes.

There are 1 outstanding health conditions. Correct these conditions before they affect array
operation.
```

Each event message includes the following information:

- Event priority: INFO, WARNING, ERROR, or FATAL (see *Event priorities* on page 14-2)
- Date and time that the event occurred
- Member on which the event occurred
- Descriptive event text

In addition, the group generates hardware alarms. An alarm is a persistent condition in an array (for example, high temperature). Alarms help you find and correct problems before they disrupt operations. An alarm always has a corresponding event. See *About hardware alarms* on page 14-2.

Event priorities

Each event has a priority. Table 14-1 lists event priorities in order of lowest (least severe) to highest (most severe). The first column lists the priorities, the second describes them.

Table 14-1: Event Priorities

Priority	Description
INFO	Informational message. Indicates an operational or transitional event that requires no action.
WARNING	Potential problem. Can become an event with Error priority if administrator intervention does not occur.
ERROR	Serious failure. Identify and correct the problem as soon as possible.
FATAL	Catastrophic failure. Identify and correct the problem immediately.

About hardware alarms

The group generates an alarm in the event of a persistent hardware condition in a group member (for example, high temperature or a failed power supply). Alarms help you discover and correct problems before they disrupt operations.

Alarms appear in the Alarms and Operations panel at the bottom of the Group Manager GUI. Click the panel header to open and close the panel.

Each alarm has a priority level, either Warning or Critical, based on the severity of the problem. An alarm always has a corresponding event. For more information, see *Monitoring alarms and operations* on page 15-10.

Event notification methods

Set up one or more event notification methods so the group notifies you when events occur.

You can configure the following event notification methods:

- **E-mail notification.** If an event occurs, the group automatically sends a message to designated e-mail addresses.

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

  See *Configuring E-Mail notification* on page 14-3.

- **E-Mail Home.** If a hardware component fails or if you update firmware, the group automatically notifies customer support.

  E-Mail Home is available to all PS Series customers, but response time and assistance is based on the validity and level of your support contract.

  See *Configuring E-Mail home* on page 14-4.

- **Remote syslog server logging.** 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. See *Configuring syslog notification* on page 14-5.

Configuring E-Mail notification

Requirement: To use e-mail notification, a group must have access to a Simple Mail Transfer Protocol (SMTP) server or e-mail relay.

1. Click `Group`, then `Group Configuration`, and then the `Notifications` tab. See the online help for information about the data fields and options.
2. In the E-Mail Event Notifications panel, select `Send E-mail to addresses`.
3. Under E-mail recipients, click `Add` and enter an e-mail address. You can enter up to five e-mail addresses to receive e-mail notifications. Then, click `OK`.
4. Under Event Priorities, select events for which you want to generate an e-mail message. See *About event messages* on page 14-1.
5. Under E-Mail Configuration Settings, click `Add` and enter an IP address for the SMTP server or e-mail relay that handles e-mail forwarding. Then, click `OK`.

   Use the *`ip_address:port`* format to specify a port number other than the default (25).

   You can enter up to three IP addresses. The group uses one SMTP server or e-mail relay at any time. The first server you specify is the default server. The group uses the other servers in the order specified, if the default server is not available. Click the up and down arrows to change the order.
6. In the `Sender in e-mail address` field, type the e-mail address that appears in the "From" field in the notification e-mail. You can use the group name at your company's e-mail address. For example: `GroupA@company.com`. When the intended recipient receives e-mail, the e-mail itself specifies which group it came from. This is helpful in multi-group environments.
7. Click `Save all changes` (Control+S).

To test the e-mail notification, select `Enable live informational messages` (you can later disable the display of informational events), log out of the group, and then log in to the group. If an e-mail recipient does not receive notification of the logout and login events, check and fix the configuration.

Changing the E-Mail notification configuration

1. Click `Group`, then `Group Configuration`, and then the `Notifications` tab.
2. In the E-Mail Event Notifications panel:
   - To disable e-mail notification, deselect `Send E-mail to addresses`.
   - To modify an e-mail address or SMTP server IP address, select the address and click `Modify`, change the address, and click `OK`.
   - To delete an e-mail address or SMTP server IP address, select the address and click `Delete`.
   - To change event priorities that result in notification, select or deselect the appropriate priorities.
3. Click `Save all changes` (Control+S).

Configuring E-Mail home

If a hardware component fails or if you update firmware, the group can automatically notify customer support through email.

Recommendation: Dell strongly recommends that you enable E-Mail Home, to expedite customer support becoming engaged in solving any problems.

E-Mail Home is available to all PS Series customers, but response time and assistance is based on the validity and level of your support contract.

Requirement: To support E-Mail Home, the group must have access to an SMTP server or e-mail relay.

1. Click `Group`, then `Group Configuration`, and then the `Notifications` tab.
2. In the E-Mail Event Notifications panel, select `Send e-mail alerts to Customer Support (E-Mail Home)`.
3. In the `Local contact e-mail` field, enter an e-mail address to receive E-Mail Home notification messages.
4. Under E-Mail Configuration Settings, click `Add` and enter an IP address for the SMTP server or e-mail relay that handles e-mail forwarding. Then, click `OK`.

   Use the *`ip_address:port`* format to specify a port number other than the default (25).

   You can enter up to three IP addresses. The group uses one SMTP server or e-mail relay at any time. The first server you specify is the default server. The group uses the other servers in the order specified, if the default server is not available. Click the up and down arrows to change the order.
5. In the `Sender in e-mail address` field, type the e-mail address that appears in the "From" field in the notification e-mail. You can use the group name at your company's e-mail address. For example: `GroupA@company.com`. When the intended recipient receives e-mail, the e-mail itself specifies which group it came from. This is helpful in multi-group environments, and reduces the chance that the e-mail server or recipient discards or rejects notifications.
6. Click `Save all changes` (Control+S).

When you first enable E-Mail Home, the group sends the local contact e-mail address a confirmation message. If you do not receive this message:

- Make sure that you specified the correct information in the Group Notifications window.
- 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.

Changing the E-Mail home configuration

1. Click `Group`, then `Group Configuration`, and then the `Notifications` tab.
2. In the Group Notifications window:
   - To disable E-Mail Home, deselect `Send e-mail alerts to Customer Support (E-Mail Home).`
   - To modify the local e-mail address, enter the new address in the `Local contact e-mail` field.

- To modify the IP address for an SMTP server, select the IP address, click `Modify`, change the address, and click `OK`.
- To delete an SMTP server, select the IP address and click `Delete`.

3. Click `Save all changes` (Control+S) to apply the changes.

Configuring syslog notification

Requirement: The syslog server must be able to store remote log files.

1. Click `Group`, then `Group Configuration`.
2. Click the `Notifications` tab. See the online help for information about the data fields and options in the Event Logs panel.
3. In the Event Logs panel, select `Send events to syslog servers`.
4. Under Syslog Servers, click `Add` and specify an IP address for the syslog server. You can specify up to three syslog servers. All the servers receive events.
5. Under Event Priorities, select the event priorities that result in syslog server notification. See *About event messages* on page 14-1.
6. Click `Save all changes` (Control+S).

Changing the syslog notification configuration

1. Click `Group`, then `Group Configuration`, and then the `Notifications` tab.
2. In the Event Logs panel:
   - To disable syslog notification, deselect `Send events to syslog servers`.
   - To modify the IP address for a syslog server, select the IP address, click `Modify`, 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.
3. Click `Save all changes` (Control+S) to apply the changes.

Enabling or disabling the display of INFO event messages

You can enable (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`, then `Group Configuration`, and then the `Notifications` tab.
2. In the Event Logs panel, select or deselect `Enable live informational messages`.
3. Click `Save all changes` (Control+S).

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.

See *Displaying and configuring SNMP access to a Group* on page 4-11 and *Configuring SNMP trap destinations* on page 14-7.

The PS Series array MIBs (Management Information Bases) contain information about SNMP traps and trap thresholds. See *Accessing PS Series array MIBs* on page 14-7.

Table 14-2 lists PS Series SNMP Traps

Table 14-2: PS Series SNMP Traps

Trap Type	Trap Names
Battery backup	eqlMemberHealthBatteryLessThan72Hours eqlMemberHealthNVRAMBatteryFailed eqlMemberHealthhighBatteryTemperature
Component	eqlMemberHealthhwComponentFailedCrit eqlMemberHealthincompatControlModule eqlMemberHealthopsPanelFailure eqlMemberHealthemmLinkFailure eqlDiskStatusChange
Fan and PSU	eqlMemberHealthFanSpeedHighThreshold eqlMemberHealthFanSpeedLowThreshold eqlMemberHealthFanTrayRemoved eqlMemberHealthBothFanTraysRemoved eqlMemberHealthPowerSupplyFailure
iSCSI	iscsiTgtLoginFailure iscsiIntrLoginFailure iscsiInstSessionFailure scsiTgtDevicesStatusChanged scsiLuStatusChanged
Link	linkUp linkDown
RAID	eqlMemberHealthRAIDSetDoubleFaulted eqlMemberHealthRAIDLostCache eqlMemberHealthRAIDSetLostBlkTableFull eqlMemberHealthRaidOrphanCache eqlMemberHealthRaidMultipleRaidSets
Security	authenticationFailure
Start	coldStart warmStart
Temperature	eqlMemberHealthTempSensorHighThreshold eqlMemberHealthTempSensorLowThreshold eqlMemberHealthlowAmbientTemp

Configuring SNMP trap destinations

You can configure network addresses to receive SNMP traps from the group.

1. Click `Group`, then `Group Configuration`.
1. Click the `SNMP` tab. See the online help for information about the data fields and options.
2. In the SNMP Traps panel, click `Add`.
3. Enter the IP address where SNMP traps are sent and click `OK`. You can specify up to five IP addresses. All the addresses receive traps.
4. Optionally, modify the SNMP trap community name. The default is `SNMP-trap`. The group uses the SNMP trap community when it sends SNMP traps. SNMP trap community names must be unique and can contain up to 63 alphanumeric characters, but no commas.
5. Click `Save all changes` (Control+S).

Changing the SNMP trap configuration

1. Click `Group`, then `Group Configuration`, and then the `SNMP` tab.
2. In the SNMP Traps panel:
   - To modify an SNMP trap destination, select the IP address, click `Modify`, change the IP address, and click `OK`.
   - To delete an SNMP trap destination, select the IP address, and click `Delete`.
   - To modify the SNMP trap community name, change the name in the `SNMP trap community name` field.
3. Click `Save all changes` (Control+S).to apply the modification.

Accessing PS Series array MIBs

PS Series array MIBs (Management Information Bases) contain information about SNMP traps and trap thresholds. The group collects MIBs in an archive. You can open this archive and examine the individual MIBs, or you can download the archive and store it on your computer.

Requirement: To use MIBs, you must install them on a management station. Not all traps apply to all array models.

If you have a support account, you can access the PS Series array MIBs from the customer support website.

1. Log in to a support account and access the `Downloads` web page.
2. Click the link for your PS Series Firmware version.
3. Click `MIBS`.

15 Group monitoring

It is best practice to regularly monitor a PS Series group, so you can address issues before service is interrupted.

About monitoring best practices

Dell recommends that you set up event notification to inform you automatically of events and operations in a group. See *Event notification methods* on page 14-2.

If you configured SNMP trap notification, you can examine the traps using an SNMP console. See *About SNMP traps* on page 14-6.

Table 15-1 describes general best practices for monitoring a group. The first column lists the monitoring condition, the second column describes it, and the third column provides a reference for more information about addressing the issue.

Table 15-1: General Monitoring Best Practices

Monitor	Description	Reference
Events of WARNING, ERROR, and FATAL severity	Investigate significant events and take the appropriate action to prevent or resolve problems.	See *Monitoring events* on page 15-2
Hardware failures	Replace failed hardware promptly. Multiple hardware failures might result in an offline member or lost data. Do not remove a failed hardware component until you are ready to replace it.	See *Monitoring alarms and operations* on page 15-10 and *Monitoring group members* on page 15-15
Degraded RAID set	If a RAID 1 or RAID 5 set is degraded, another disk drive failure in the RAID set might result in lost data. If a RAID 6 set is degraded, one or two additional drive failures might result in lost data. Immediately replace failed drives.	See *Monitoring a specific member* on page 15-16
Offline volumes	An offline volume might indicate a problem. For example, a member might be offline.	See *Monitoring volumes and snapshots* on page 15-25
Low pool space	Do not let pool space fall below a recommended value.	See *Monitoring storage pool free space* on page 15-15
Low free volume space	If there is insufficient free volume space, writes to the volume fail. This can affect application performance. If a write to a volume exceeds the reported size, it fails. This can affect applications that use the volume. If a thin-provisioned volume reaches its maximum in-use space limit (and the limit is less than 100%), the group sets the volume offline.	See *Monitoring volumes, collections, and snapshots* on page 15-24
Incomplete pool move operations	Moving a volume or a member from one pool to another can take a long time. Monitor the progression of the move operation and make sure that it completes.	See *Monitoring group operations* on page 15-14

Getting started with group monitoring

All monitoring information is available by clicking `Monitoring` in the lower-left panel.

Monitoring events

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 you create a volume).

Event messages help you monitor group operations and correct problems before they disrupt operations. See *About event messages* on page 14-1.

To display events, click `Monitoring` and then `Events`.

Table 15-2 shows the filtering options available in the Events panel.

Table 15-2: Events Panel

Option	Description	Shortcut	User Actions
View	Filters the list by event type with to include all events, warnings and errors, or critical errors.	Alt+V	See *About monitoring best practices* on page 15-1 See *Monitoring alarms* on page 15-10
More	The events list displays events for a specified period of time. Select this option to display events from a previous time period.	Alt+M	See *About monitoring best practices* on page 15-1
Acknowledge all	Acknowledges all new events.	None	See *About monitoring best practices* on page 15-1
Clear event list	Deletes all events in the list.	None	See *About monitoring best practices* on page 15-1
View details	Provides detailed information about the selected event.	None	See *About monitoring best practices* on page 15-1
Hide details	Hides details for the selected event.	None	See *About monitoring best practices* on page 15-1

From the Group Events window, you can:

- Display all events or events of a specific priority. Pull down the `View` menu and select the events you want to display.
- Retrieve previous events. To retrieve the most recent 100 events, click the `More` icon (). 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 the `More` icon.
- Show or hide details, by doing any of the following:
  - Move the pointer over an event. A pop-up window appears, 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 `View details` () or `Hide details` icons (). The event details panel opens at the bottom of the events list.

See the online help for information about the data fields and options.

Accessing the event log file on a remote computer

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

See *Configuring syslog notification* on page 14-5.

Accessing events sent to an E-Mail address

If you configured e-mail notification of events, you can view the events by logging into one of the e-mail accounts that you configured for notification.

See *Configuring E-Mail notification* on page 14-3.

Monitoring administrative sessions

To monitor administrative statistics, click `Monitoring` and then `Administrative Sessions`. The Administrative Sessions window appears, containing the following panels:

- Active administration sessions – Provides information about the login:
  - Account name, session type, and login time
  - IP of local group and client computer
- Administrative login history – Provides information about account usage:
  - Account name, authentication and type
  - Pool access and last login

Double-click the name of a user to open the Modify Administration Account wizard.

See the online help for information about the data fields and options.

Monitoring iSCSI connections

To monitor iSCSI connection statistics for all the targets (volumes and snapshots) in the group, click `Monitoring` and then `iSCSI Connections`. The iSCSI Connections panel provides the following information:

- Initiator address and target name.
- Connection time and connection ethernet interface
- Data transfer volume

Check for multiple initiators writing to the same target. This can cause volume corruption if not handled correctly by the servers.

Note: You can sort the table in the GUI by clicking column headings. By default, the table is sorted by Initiator address.

See the online help for information about the data fields and options.

Monitoring snapshot schedules

To monitor snapshot schedules, click `Monitoring` and then `Snapshot Schedules`. The Snapshot Schedules panel provides the following information:

- Name of schedule and the associated volume or collection
- Type of object created (such as snapshot)
- Time and data parameters and schedule run status

To display more detail about a schedule, move the pointer over a schedule entry in the panel. A pop-up window appears, showing additional information such as the:

- 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 the Modify icon in the Actions column. The `Modify Schedule` dialog box opens. Make the changes, then click `OK`. You can change the:
  - Schedule name
  - Run date and time
  - Number of snapshots to keep
- To disable or enable a schedule, select it and click the flag icon (enable or disable). In the confirmation dialog box, click `Yes` to disable the schedule.
- To delete a schedule, select it and click the X icon (delete). In the confirmation dialog box, click `Yes` to delete the schedule.

See the online help for information about the data fields and options.

Monitoring replication schedules

To monitor replication schedules, click `Monitoring` and then `Replication Schedules`. The Replication Schedules panel provides the following information:

- Name of schedule and the associated volume or collection
- Type of object created (such as snapshot)

- Time and data parameters and schedule run status

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 such as the:

- 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 the Modify icon in the Actions column. The `Modify Schedule` dialog box opens. Make the changes, then click `OK`. You can change the:
  - Schedule name
  - Run date and time
  - Number of replicas to keep
- To disable or enable a schedule, select it and click the flag icon (enable or disable). In the confirmation dialog box, click `Yes` to disable the schedule.
- To delete a schedule, select it and click the X icon (delete). In the confirmation dialog box, click `Yes` to delete the schedule.

See the online help for information about the data fields and options.

Monitoring replication

If you are replicating volume data, you should monitor replication operations to ensure 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)

See the online help for information about the data fields and options.

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.

For details about displaying information about replication partners, see *Monitoring replication partners* on page 15-9.

Table 15-3 describes some best practices for monitoring replication between groups. The first column lists the monitoring condition, the second column describes it, and the third column provides a reference for more information about addressing the issue.

Table 15-3: Replication Monitoring Best Practices

Monitor	Description	Reference
Incomplete replication operations	If a replication operation fails to complete, you might need to increase replication space.	See *Monitoring replication* on page 15-5
Incomplete manual transfer operations or failback operations	Some operations require multiple tasks that administrators must complete. Make sure you complete all multi-task operations.	See *Monitoring replication* on page 15-5
Low free delegated space	If delegated space is low, replica reserve space might not be able to increase automatically to store new replicas.	See *Monitoring a specific partner* on page 15-10 and *Monitoring inbound replication* on page 15-8
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.	See *Monitoring inbound replication* on page 15-8

Monitoring outbound replication

Outbound replication transfers volume data from the current group to a replication partner.

To monitor outbound replication, click `Monitoring` and then `Outbound Replication`. The Outbound Replication panel provides the following information:

- Volume and partner name
- Transfer status, start time, and amount of data transferred

See the online help for information about the data fields and options.

To see more detail about a volume in the list, move the pointer over the volume name. A pop-up window appears, listing the following:

- Volume name
- Storage pool
- Thin provisioning status (enabled or disabled)
- Reported size, volume reserve, and snapshot reserve
- Current and requested status (online or offline)
- Access type (read-write permission and whether shared for cluster access)
- Replication partner
- Description, if any

You can take the following actions:

- Click a volume name to navigate to the Volumes Status window.
- Double-click a row to navigate to the Outbound Replicas window.

Table 15-4 shows the outbound replication operation status. The first column lists the status values, the second column provides descriptions, and the third column provides solutions.

Table 15-4: Outbound Replication Operation Status

Status	Description	Solution
`authfailure`	Failed authentication between the replication partners.	Make sure you configured the partners with the correct passwords.
`cancelling`	Administrator cancelled the replication operation.	None needed; informational
`completed`	Most recent replication is complete.	None needed; informational
`disabled`	Administrator disabled replication for the volume.	None needed; informational
`failed`	Volume replication failed.	Examine the event log for information about the failure and correct the problem.
`inprogress`	Volume replication is in progress.	None needed; informational
`manual-transfer-in-progress`	Manual transfer is in progress.	None needed; informational
`partner-down`	Volume replication cannot continue because the partner is unavailable.	Make sure the network link between the partners is configured correctly and functioning.
`partner-paused-inbound`	Partner administrator paused inbound replication.	None needed; informational
`partner-paused-outbound`	Partner administrator paused outbound replication.	None needed; informational
`pause-max-snaps-reached`	Replication paused because the secondary group contains the maximum number of replicas or snapshots for a group.	Delete replicas or snapshots on the secondary group.
`paused`	Administrator paused replication to the partner.	None needed; informational
`paused-remote-reserve-low`	Replication paused because of insufficient replica reserve.	On the primary group, increase the replica reserve percentage. An event message specifies the amount to which the replica reserve must be increased for the replication to complete. Replication continues automatically.
`paused-remote-resize-failed`	Replica reserve resize operation failed due to insufficient delegated space.	On the secondary group, increase the space delegated to the primary group.
`ready`	Replication between the partners is correctly configured and is ready to process new replication operations.	None needed; informational.
`remote-disallow-downgrade-not-set`	Replication paused because the secondary group did not disallow firmware downgrades.	On the secondary group, disallow downgrades.
`remote-partner-needs-upgrade`	Replication paused because the secondary group is running incompatible firmware.	Upgrade secondary group firmware to the same firmware as the primary group. Replication resumes automatically.

Table 15-4: Outbound Replication Operation Status (Continued)

Status	Description	Solution
`remote-replicaset-is-recovery-volume`	The replica set on the partner has been promoted to a recovery volume.	Demote the recovery volume on the partner and retry the replication.
`waiting`	Replication data transfer did not start because the group cannot create more iSCSI sessions.	In most cases, the problem resolves itself, and replication continues automatically.

Monitoring inbound replication

Inbound replication stores volume data on the current group, in replica sets, from a replication partner.

To monitor inbound replication, click `Monitoring` and then `Inbound Replication`. The Inbound Replication panel provides the following information:

- Volume and partner name
- Replication start time and status
- Amount of data transferred

See the online help for information about the data fields and options.

Table 15-5 shows the inbound replication operation status, its description, and how to solve any issues.

Table 15-5: Inbound Replication Operation

Status	Description	Solution
`authfailure`	Authentication between the partners failed.	Make sure you configured the partners with the correct passwords.
`farenddown`	No communication with the primary group for one hour or more.	Correct the network problems or make sure the primary group is available.
`inprogress`	Replication is in progress.	None needed; informational.
`ready`	Replication is correctly configured on the volume.	None needed; informational.
`stopped`	Administrator paused replication.	None needed; informational.
`unmanaged`	Delegated space capacity on the secondary group that is no longer accessible from the primary group. This can happen when the original volume is no longer available or no longer configured for replication.	If these replicas are no longer needed, you can delete them or permanently promote them to a volume.
`waiting`	Replication data transfer did not start because the group cannot create more iSCSI sessions.	In most cases, the problem resolves itself, and replication continues automatically.

To see more detail about a replica set in the list, move the pointer over the replica set name. A pop-up window appears, listing the:

- Replica set name
- Storage pool where the replica set is stored on the current group
- Thin provisioning status (enabled or disabled)

- Maximum reserve, snapshot reserve, and free space
- How the latest replication occurred (over the network or through Manual Transfer Replication)
- Whether a failback snapshot is enabled
- Current and operational status (online or offline)
- Number of replicas in the replica set
- Description, if any

Click a replica set name or double-click a row to navigate to the Inbound Replicas window for that partner, with the replica set selected and its detail shown.

Monitoring outbound replication history

To see replication history, click `Monitoring` and then `Replication History`. The Outbound Replication History panel provides the following information:

- Volume and partner name
- Replication start time, duration, and status
- Amount of data transferred and the transfer speed

See the online help for information about the data fields and options.

Table 15-6 describes replica status and descriptions.

Table 15-6: Replica Status

Status	Description
`complete`	Replication operation is complete, and all the data is on the secondary group.
`incomplete` or `in progress`	Replication operation is not complete. A replication operation cannot start if another operation replication for the same volume is in progress.

Periodically 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.

In the Transferred column, 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 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 (replication from other groups to this group).

To display a list of all the replication partners for a group, click `Replication` and then `Replication Partners`.

The Delegated Space panel shows all the delegated space for all partners, and how much free space is available. The Replication panel shows all the replications, and their direction, between this group and all configured partners.

You should monitor the usage of delegated space. If free delegated space is not available, replica reserve cannot increase automatically.

Monitoring a specific partner

To display details about a specific partner, click `Replication` and then the replication partner.

The General Partner Information panel shows the partner name, IP address, and contact information.

The 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.

The Replication Progress panel shows any in-process replication activity between the groups.

For each partner, you can display details about the following:

- Inbound replica collections
- Inbound replicas
- Outbound replica collections
- Outbound replicas

Monitoring 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. See *About hardware alarms* on page 14-2.

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. See *Displaying warning alarms* on page 15-12.

- Critical – Serious problem that can cause damage to the array or data loss. See *Displaying critical alarms* on page 15-11.

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.

Displaying critical alarms

Open the Alarms and Operations panel and click the `Critical` tab to display Critical alarms.

Table 15-7 shows the data fields available in the Critical tab.

Table 15-7: Alarms and Operations Panel - Critical Tab

Column	Description	User Actions
Severity	Severity of the alarm.	*Monitoring alarms and operations* on page 15-10 *Accessing the alarms panel* on page 3-4
Object	Object to which the alarm applies.	*Monitoring alarms and operations* on page 15-10 *Accessing the alarms panel* on page 3-4
Condition	Condition that triggered the alarm.	*Monitoring alarms and operations* on page 15-10 *Accessing the alarms panel* on page 3-4

Critical alarms appear on the `Critical` tab of the Alarms panel. 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.

Note: Critical alarms correspond to ERROR events.

Critical alarms include:

- 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).

Displaying warning alarms

Open the Alarms and Operations panel and click the `Warnings` tab to display Warning alarms.

Table 15-8 shows the data fields available in the Warning tab.

Table 15-8: Alarms and Operations Panel - Warning Tab

Column	Description	User Actions
Severity	Severity of the alarm.	*Monitoring alarms and operations* on page 15-10
Object	Object to which the alarm applies.	*Monitoring alarms and operations* on page 15-10
Condition	Condition that triggered the alarm.	*Monitoring alarms and operations* on page 15-10

Warning alarms appear on the `Warnings` tab of the Alarms panel. A warning alarm indicates a condition that decreases performance or can become critical if you do not correct it.

Note: Warning alarms correspond to WARNING events.

Warning alarms include:

- 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.

Monitoring actions

Open the Alarms and Operations panel and click the `Actions` tab to display any steps that you must complete. See *Monitoring alarms and operations* on page 15-10.

Table 15-9 shows the data fields available in the Actions tab.

Table 15-9: Alarms and Operations Panel - Actions Tab

Column	Description	User Actions
Severity	Severity of the action.	*Monitoring group operations* on page 15-14
Object	Object to which the action applies.	*Monitoring group operations* on page 15-14
Condition	Action taking place.	*Monitoring group operations* on page 15-14

Some complex operations, such as manual transfer replication, require administrators to perform multiple tasks.

To display incomplete tasks, open the Alarms and Operations panel and click the `Actions` tab.

If a multi-task operation is in progress, the GUI displays the incomplete tasks. Make sure you complete all multi-task operations.

Monitoring group operations

Open the Alarms and Operations panel and click the `Group Operations` tab to display group management operations (for example, moving a member to another pool) and actions you might need to take.

Table 15-10 shows the data fields available in the Group Operations tab.

Table 15-10: Alarms and Operations - Group Operations Tab

Column	Description	User Actions
Started	Date and time the operation started.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4
Object	Object on which the operation is performed.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4
Operation	Type of operation performed.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4
Status	Status of the operation.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4
Progress	Progress of the operation.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4
Actions	Contains an option for canceling in-progress operations.	*About event messages* on page 14-1 *Accessing the alarms panel* on page 3-4

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.

Monitoring failback operations

Open the Alarms and Operations panel and click the `Failback Operations` tab to display failback operations and any actions you might need to take.

Table 15-11 shows the data fields available in the Failback Operations tab.

Table 15-11: Alarms and Operations - Failback Operations Tab

Column	Description	User Actions
Name	Name of the recovery volume participating in the failback operation. Expand the volume to display the status of the individual tasks in the operation.	*Monitoring replication partners* on page 15-9 *Monitoring failback operations* on page 15-14
Executing On	Partner on which the operation is taking place.	*Monitoring replication partners* on page 15-9 *Monitoring failback operations* on page 15-14

Table 15-11: Alarms and Operations - Failback Operations Tab (Continued)

Column	Description	User Actions
Status	Status of the operation.	*Monitoring replication partners* on page 15-9 *Monitoring failback operations* on page 15-14
Started	Time and date the operation started.	*Monitoring replication partners* on page 15-9 *Monitoring failback operations* on page 15-14

Monitoring 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.

1. Click `Group` and then `Storage Pools`. The Storage Pool Summary window appears.
2. Check the free pool space value in the Storage Pools panel. Dell recommends that free pool space does not fall below the following, whichever is smaller:
   - 5% of the total pool space
   - 100 GB 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. See *Moving a volume to a pool* on page 7-5.

You can expand pool capacity by moving a member to the pool.

To sort the table, by click a column heading. The table is sorted by Volume Template by default.

Monitoring group members

Member hardware problems typically cause event messages and alarms. Monitor the member hardware and replace any failed components immediately.

1. Click `Group` and then `Members`. The `Group Disk Space` panel shows the total amount of free space in the group and in each pool (if applicable).

   The `Group Members` panel lists all members, the pool to which each member belongs, the capacity and amount of free space, RAID policy, number of disks, status, PS Series Firmware version (should be the same for all members), and the number of iSCSI connections to each member. (This indicates the number of volumes or snapshots with data on that member that are connected to an initiator. Nothing connects directly to a member.)
2. Check the following:
   - Member status – Table 15-13 describes 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).

Monitoring a specific member

Click `Group`, expand `Members`, then select the member name, and then click the `Status tab.`

Displaying general member information

In the General Member Information panel, check the RAID status. Table 15-12 describes the RAID status values and provides possible solutions where appropriate.

Table 15-12 shows RAID status for a member. The first column lists the status values, the second column provides descriptions, and the third column provides solutions.

Table 15-12: RAID Status

Status	Description	Solution
`catastrophicLoss`	Disk array lost group metadata or user data. The array does not initialize.	Contact your support provider.
`degraded`	A RAID set is in a degraded state. If the member RAID level is RAID 5, RAID 50, or RAID 6, performance might be impaired.	Identify and replace any failed drives.
`expanding`	Disk array is expanding (for example, because you installed additional disk drives or changed the member RAID policy).	None needed; informational.
`failed`	Multiple disk drive failures occurred in the same RAID set. The member is set offline.	Contact your support provider.
`ok`	Disk array initialization is complete. Performance is normal.	None needed; informational.
`reconstructing`	Disk array is reconstructing data on a drive (for example, because a drive failed, and a spare is replacing it). During reconstruction, performance might decrease. After reconstruction, performance returns to normal, unless a RAID set is degraded.	Identify and replace any failed drives.
`verifying`	Disk array is initializing (for example, because you set the member RAID policy).	None needed; informational.

Displaying member health status

1. Click `Group`, then expand `Members`, then select the member name, and then click the `Status tab.`
2. In the Member Health Status panel:
   - Click `Front view` to display the front panel of the array.
   - Click `Rear view` to display the back panel of the array, including the control modules and the power supply and cooling modules. The front and rear views shown in your GUI depend on the array model of the group member.
   - Click `Inside view` (not available on all array models) to display the interior disk drive slots.
   - Click `View alarms` to display all the alarms for the member.

A red X over a hardware component indicates uninstalled or unconfigured hardware. A warning or error status symbol in the array graphic indicates a failed or failing component. Move the pointer over a component to show status details.

Table 15-13 shows member status, description, and how to solve any issues.

Table 15-13: Member Status

Status	Description	Solution
`unconfigured`	You did not select a RAID policy for the member.	None needed; informational.
`initializing`	Member is initializing according to the selected RAID policy.	None needed; informational.
`online`	Array is a functioning member of the group.	None needed; informational. A member can experience a failure but still be online.
`offline`	Member is unavailable, failed, or power was removed.	Identify and correct the problem.
`vacating-in-progress`	Member is moving data to the remaining pool members before it is removed from the group.	None needed; informational. This can be a long operation, based on the amount of data that must be moved to the other pool members.
`vacated`	Member has successfully moved its data to the other pool members before it is removed from the group.	None needed; informational.

Displaying member space

Click `Group`, then expand `Members`, then select the member name, and then click the `Status tab`.

The Member Space panel shows the total amount of usable space on the member, how much space is used by volumes, snapshots, and replicas, and the amount of free space, numerically and in a graphic.

Using LEDs to identify a member

If a hardware failure occurs in a member, the group generates an alarm which causes the member LED to light.

In addition, to help you identify a member, you can make the fan tray LED and the control module ERR LED on the member chassis flash.

- To make a member's LED flash, click `Group`, then expand `Members`, then select the member name, and then click `Start LEDs flashing`.
- To stop flashing a member's LED, click `Group`, then expand `Members`, then select the member name, and then click `Stop LEDs flashing`

Warning: Never turn off power to a group member unless the member has been cleanly shut down. See *Shutting down a member* on page 6-14.

Monitoring 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 card.

To display the member enclosure information, click `Group`, then expand `Members`, then select the member name, and then click the `Enclosure` tab.

Monitoring power supplies

A member has two or three power supplies. Most PS Series arrays use power supplies that have integrated cooling modules.

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 Maintenance* 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.

Table 15-14 shows power supply status and possible solutions. The first column lists the status values, the second column provides descriptions, and the third column provides solutions.

Table 15-14: Power Supply Status

Status	Description	Solution
`OK`	Array is receiving power from the power supply.	None needed; informational.
`no-power`	Power supply is not installed or not connected to a power source, or the power supply is not turned on (not all power supply models).	Keep all power supplies installed and connected to a power source. If the power supply has a power switch, make sure the power switch is on.
`failed`	Power supply failure.	See your PS Series support provider for information about replacing the power supply.

Monitoring cooling 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, feel 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.

The Cooling Fans panel shows the status of the fans on the cooling modules. Table 15-15 shows the cooling fan status. These status values apply to array models with combination power supply and cooling modules. The first column lists the status values, the second column provides descriptions, and the third provides solutions.

Table 15-15: Cooling Fan Status

Status	Description	Solution
`fan-present`	Cooling modules and fan are functioning.	None needed; informational.
`fan-not-present`	Cooling module or fan failed, or the cooling module is not installed, not turned on, or not connected to a power source.	Install a functioning cooling module, make sure the cooling module is connected to a power source, or turn on the cooling module (not available on all cooling module models). See your PS Series support provider for information about replacing a failed cooling module.

The Temperature Sensors panel shows the current temperature for the various array controllers and processors, in addition to the normal temperature range. Table 15-16 describes the array temperature status, descriptions, and how to solve any issues.

Table 15-16: Array Temperature Status

Status	Description	Solution
`normal`	Temperature is within normal range.	None needed; informational.
`warning`	Temperature is outside normal range, but within limits.	Check that all fans are working properly. Monitor the temperature carefully.
`critical`	Temperature is outside operating limits.	Check that all fans are working properly. Make sure the air conditioning system is working correctly, and make sure there is air flow around the array. If a processor temperature stays high, replace the control module. See your PS Series support provider for information about replacing a failed cooling module.

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
```

Monitoring 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.

Table 15-17 shows channel card status, descriptions, and how to solve any issues.

Table 15-17: 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.

For information about replacing channel cards, see the *Hardware Maintenance* manual for your array model or contact your PS Series support provider.

Monitoring the EIP card

Some array models include an Enclosure Interface Processor (EIP) card. An array continues to operate if the EIP card fails. You can replace the failed EIP card with no impact on group operation.

In the Member Enclosure window, the EIP card panel shows the EIP card status.

Table 15-18 describes EIP card status, descriptions, and how to solve any issues.

Table 15-18: EIP Card Status

Status	Description	Solution
`good`	EIP card is functioning normally.	None needed; informational.
`failed`	EIP card failure.	Contact your PS Series support provider for information about replacing an EIP card.
`not-present`	EIP card is missing or status is unavailable.	Contact your PS Series support provider for information about installing or replacing an EIP card.

For information about replacing the EIP card, see the *Hardware Maintenance* manual for your array model or contact your PS Series support provider.

Monitoring 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). Upon 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. This is 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 Maintenance* manual for your array model or contact your PS Series support provider.

To display control module information, click `Group`, then expand `Members`, then select the member name, and then click the `Controllers` tab.

Each Control Module Slot panel shows the following information:

- Status. See Table 15-19.
- Boot time.

- Cache battery status and NVRAM battery status. See Table 15-20 and Table 15-21 for descriptions of battery status and possible solutions where appropriate.
- Model number.
- Boot ROM version.
- PS Series firmware version.

An empty slot means that a control module is not installed or has failed.

For information about replacing a control module, see the *Hardware Maintenance* 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. See *About write cache operations* on page 6-11.

Control module status

Table 15-19 describes the control module status, descriptions, and how to solve any issues.

Table 15-19: 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 15-20 describes the control module cache battery status, descriptions, and how to solve any issues.

Table 15-20: Cache Battery Status

Status	Description	Solution
`ok`	Battery is fully charged.	None needed; informational.
`failed`	Battery failure.	Contact your service provider for information about replacing batteries.
`missing battery`	Battery is missing.	Contact your service provider for information about replacing batteries.
`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 information about replacing batteries.
`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 information about replacing batteries.
`good battery, is charging`	Battery is charging but has enough charge for normal operation.	None needed; informational.

NVRAM battery status

Table 15-21 describes the control module NVRAM coin cell battery status, descriptions, and how to solve any issues. Not every array has an NVRAM battery.

Table 15-21: NVRAM Battery Status

Status	Description	Solution
`good`	Battery installed and fully charged.	None needed; informational.
`bad`	Battery failure.	Contact your PS Series service provider for information about replacing batteries.
`not-present`	Battery is not installed.	Contact your PS Series service provider for information about replacing batteries.
`unknown`	Battery status is not known.	Contact your PS Series service provider for information about replacing batteries.

Monitoring disk drives

Make sure you detect and replace failed disk drives as soon as possible. Although spare disks and RAID protect data against disk failures, multiple disk failures might put data in jeopardy.

To display the disk drive information, click `Group`, expand `Members`, then select the member name, and then click the `Disks` tab.

The Disk Array Summary panel shows the disk 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 disk, including the slot, type, model and revision, size, status, and errors. Closely monitor drives with errors.

Disk drive status

Table 15-22 shows disk drive status, descriptions, and how to solve any issues.

Table 15-22: Disk Drive Status

Status	Description	Solution
`too-small`	Disk drive is smaller than other drives in the member. The drive cannot be used in the member.	Replace the drive with a drive that has the same size or a greater size than the installed drives.
`failed`	Disk drive failure.	See your PS Series support provider for information about replacing failed disk drives.
`foreign`	Disk 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 disk drive.	See your support provider. To use the drive, click `history-of-failure` and agree to use the disk.
`offline`	Indicates that the disk drive does not fall into the other status categories.	See your PS Series support provider.
`online`	Disk drive is functioning.	None needed; informational.

Table 15-22: Disk Drive Status (Continued)

Status	Description	Solution
`spare`	Disk drive is a spare drive.	None needed; informational.
`unsupported-version`	Disk drive cannot use the firmware running on the member.	See your PS Series support provider.

Warning: A disk drive failure in a RAID 5 or RAID 10 set that is degraded might result in data loss.

When a drive in a RAID set fails, a member behaves as follows:

- **If a spare disk 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 disk 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.
- **If a spare disk 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 disk, a member behaves as follows:

- **If a spare disk 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 volumes 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 disk drive. The member automatically copies the data on the failing disk drive to a spare disk 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 disk drive completely fails during the operation, data is reconstructed on the spare using parity data, as usual.

Replace any failed disks immediately. For information about replacing disk drives, see the *Hardware Maintenance* manual for your array model or contact your PS Series support provider.

Monitoring 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. This restores the group to normal full operation, including network communication.

To display the network information, click `Group`, expand `Members`, then select the member name, and then click the `Network` tab.

The Status of Network Interfaces panel shows the following information:

- Operational status – This is the current status of the network interface and can be:
  - `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.
- Requested status – This status is set by administrative action:
  - `enabled` – Configured and serving I/O.
  - `disabled` – Not serving I/O. Might be configured.

  If the operational status is `down`, but the requested status is `enabled`, identify and correct the error.

  To protect against network interface or port failure, connect multiple network interfaces on both control modules to the network.
- Speed – Make sure that the interface speed is adequate.
- MTU size – The path MTU size depends on the iSCSI initiator setting.
- Packet errors – A few packet errors are not usually a problem. If a large number of packet errors occur, network problem or a network interface or port failure might exist. Identify and correct the problem.

The IP Configuration panel shows each interface and its IP address, netmask, MAC address and description, if any.

Monitoring iSCSI connections to a member

To display all connections to a member, click `Group`, expand `Members`, select the member name, and then 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 can cause target corruption if not handled correctly by the servers.

Monitoring volumes, collections, and snapshots

You can monitor volume and snapshot status, volume space usage, volume free space, and snapshot reserve.

Monitoring volumes and snapshots

Make sure volumes and snapshots are online or offline. Make sure that thin-provision volumes are not about to run out of volume reserve. Also check free volume space.

To display information for all group volumes, click `Volumes` in the lower-left panel and then click `Volumes` in the far-left panel. The Volume Summary window appears.

To display detailed information about a specific volume, expand `Volumes` and select the volume name. The Volume Status window appears. To display snapshot information, expand a volume name and select a snapshot timestamp. The Snapshot Status window appears.

Monitor volume and snapshot space as follows:

- 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. See *Increasing the reported size of a volume* on page 10-5.

- 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% and a write exceeds this value. If the maximum in-use space is set to 100%, and a write exceeds this value, the write fails, but the volume remains online.

  To increase free volume space, increase the reported volume size. See *Increasing the reported size of a volume* on page 10-5. You can also increase free volume space by increasing the value of the maximum in-use space value up to 100%.

- 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 free space for new snapshots, or the volume and its snapshots are set offline.

  To increase the snapshot reserve or change the space recovery setting, see *About snapshot reserve settings* on page 11-2.

Each volume and snapshot has two status values:

- Requested status – Administrator-applied setting for the volume or snapshot. For example, an administrator can set a volume online or offline. See Table 15-23.
- Current status – Actual status of the volume or snapshot, regardless of the requested status. See Table 15-24.

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% and a write exceeds this value.

In the Volume Summary window, check the current volume status in the Volumes panel. The requested status appears when you move the pointer over the volume in the Volumes panel. If the requested status and current status are not the same, investigate further.

In the Volume Snapshots window, check the current snapshot status in the Snapshots panel. The requested status appears when you move the pointer over the snapshot in the Snapshots panel. If the requested status and current status are not the same, investigate further.

Volume and snapshot requested status

Table 15-23 shows the possible values for the requested status for a volume or snapshot, descriptions, and how to solve any issues.

Table 15-23: Requested Volume and Snapshot Status

Status	Description	Solution
`online`	Administrator set the volume or snapshot online.	None needed; informational.
`offline`	Administrator set the volume or snapshot offline. Computers cannot access an offline volume or snapshot.	None needed; informational.
`online-lost-cached-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 re-written, no error occurs, and the block no longer shows a status of lost.	None needed; informational.

Volume and snapshot current status

Table 15-24 shows the possible values for the current status for a volume or snapshot, descriptions, and how to solve any issues.

Table 15-24: 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.
`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. See *About snapshot reserve settings* on page 11-2.

Table 15-24: Current Volume and Snapshot Status (Continued)

Status	Description	Solution
`offline-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. See *Modifying the thin provisioning space settings* on page 10-4 or *Increasing the reported size of a volume* on page 10-5.
`offline-missing-pages`	A volume or snapshot was set offline because some volume data cannot be found. This is a serious condition.	Contact your PS Series support provider.
`offline-nospace-auto-grow`	The thin-provisioned volume and its snapshots were set offline because there was not enough free pool space 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. See *Moving a member to a pool* on page 7-4 and *Moving a volume to a pool* on page 7-5.
`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-cached-blocks`	Volume or snapshot was automatically set offline because blocks were lost. Computers cannot access the volume or snapshot.	Click the status link and select how to manage the lost blocks. See *Managing a volume or snapshot with lost blocks* on page 10-17 for more information.

Using the Performance Monitor

Use the Performance Monitor to show performance statistics for the drives or control modules in a member. The Performance Monitor collects statistical data every second.

You can start the Performance Monitor from the Tools menu, or from the `Members` panel in the Group Manager GUI.

Starting Performance Monitor from the tools menu

1. Open the Tools menu and click Performance monitor. The Performance Monitor GUI starts.
2. Click `Add statistics`.
3. In the `Select Statistics` dialog box:
   a. Expand `Members`
   b. Expand a specific member
   c. Expand a component or statistics category
   d. Select a statistic to display
   e. Click `OK`

The data is displayed in the Performance Monitor main window (Figure 15-1).

Starting Performance Monitor from the Group Manager GUI

1. From the `Members` list, select a member name.
2. On the Disks or Network tab, click a row for a drive or an Ethernet port.
3. In the Statistics Activity panel, select the statistic to display.

The Performance Monitor window opens and displays the selected data (Figure 15-1).

Recommendation: While it is possible to open several Performance Monitor windows at the same time, Dell recommends running only one to avoid possible unexpected results.

Figure 15-1: Performance Monitor

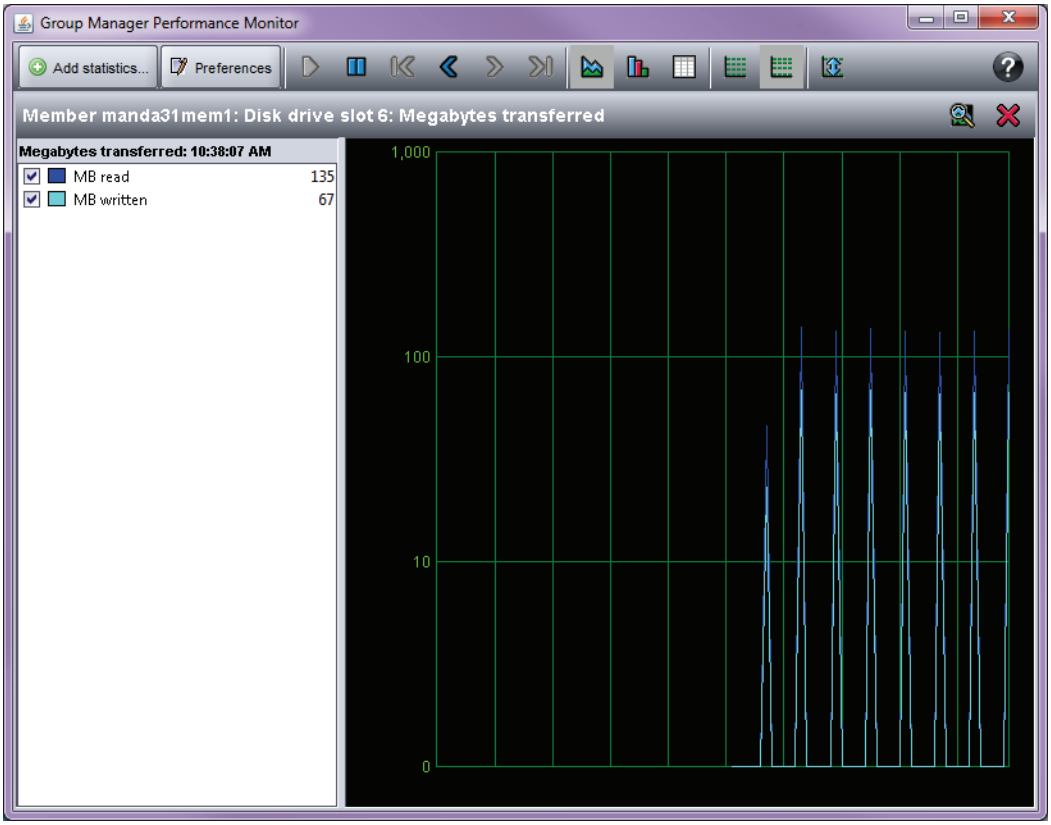

Using the Performance Monitor

Table 15-25 shows the operation icons in the Performance Monitor window.

Table 15-25: Performance Monitor Operations

Click 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).

Adding, changing, or removing statistics

You can display up to four sets of statistics in the Performance Monitor.

To add more statistics:

1. Click `Add statistics`.
2. In the `Select Statistics` dialog box:
    a. Expand `Members`
    b. Expand a specific member
    c. Expand a component or statistics category
    d. Select a statistic to display
    e. Click `OK`
3. To add more sets of statistics (up to four), repeat steps 1 and 2.

In the header of each statistics panel, the following icons perform additional operations:

- Click  to select a different data set to display in that panel
- Click  to close (delete) that panel. 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).

Changing how data is displayed

Table 15-26 shows the icons you use to change the data display. You can view the data as a chart, histogram (bar graph), or as a data table.

Table 15-26: 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 or to fit)
	Histogram (bar graph)	 - Displays data on a linear scale  - Displays data on a logarithmic scale  - Resizes the display (scales in or out or to fit)
	Data table	None

Note: The display mode you select applies to all panels currently open in the Performance Monitor window.

Displaying 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 Figure 15-2). In this example, at the peak of the graph, the group is processing 303 output requests.

Figure 15-2: Performance Monitor - Select Data Point

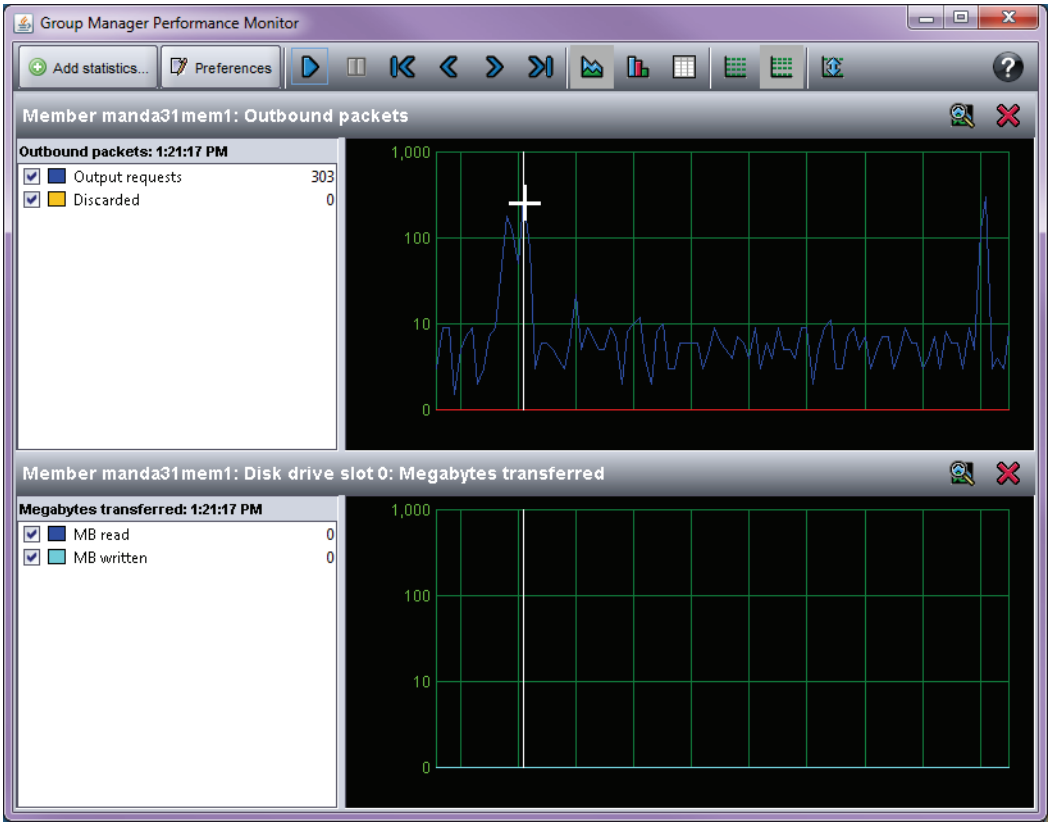

Customizing the Performance Monitor

Within the Performance Monitor window, you can change the following:

- colors used in graphs
- length of time between which data points are collected
- number of data points to save

Changing the display colors

On the left side of the window within each statistics panel, you can change the colors used in the display. Click the colored box (for example, ■) to open the Select a color dialog box (Figure 15-3). You can choose from a predefined swatch panel, or open the HSB or RGB tabs to specify a custom color value. Click OK when done.

Figure 15-3: Performance Monitor- Select a Color Dialog Box

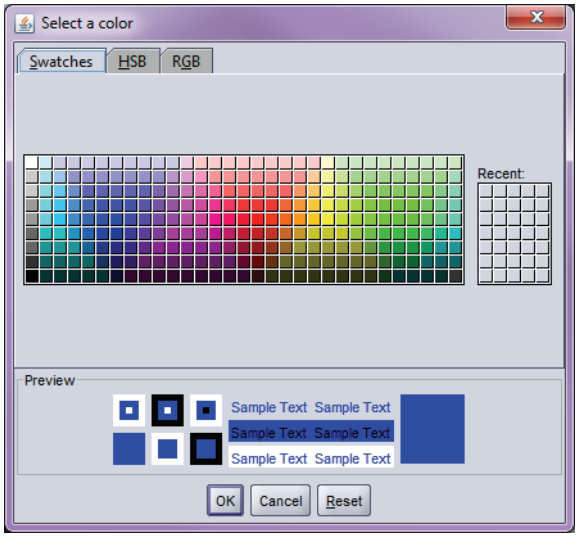

Changing the data collection values

Click `Preferences` to change:

- The time interval between data points. The default is 1 second, and the maximum is 60 seconds. You can choose from a list of predefined intervals (1, 5, 10, 30, and 60) or enter any integer value between 1 and 60; for example, 45.
- The number of data points to save. The default is 100 and the maximum is 1000. You can choose from a list of predefined values (100, 250, 500, or 1000) or enter any integer value between 100 and 1000; for example, 300.

  As you choose different values for the interval and number of data points to save, the dialog box shows you the total amount of time 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.

Click `OK` to save your changes. The changes take effect immediately; you do not have to stop and restart polling.

Additional monitoring tools

To enhance your view of PS Series group operation, you can use additional tools:

- SAN HeadQuarters – Enables you to monitor multiple PS Series groups from a single graphical interface. It gathers and formats performance data and other vital group information. Analyzing the data might help you improve performance and more effectively allocate group resources.

  Visit the customer support website for more information about SAN HeadQuarters.
- Multiple Router Traffic Grapher (MRTG) – A supplement to the Performance Monitor, enables you to monitor the I/O activity, latency, and throughput of volumes and members.

  For more information on using MRTG, see the *CLI Reference* manual.

Contacting customer support

The Customer Support website contains downloads for firmware updates, documentation, and other services. You can also create and log into your customer support account to report a problem and receive direct technical support.

To launch the EqualLogic Customer Support website from the GUI, click `Tools` and then `Customer Support`.

For more information about creating or logging into a support account, and receiving technical support, see *Technical Support and Customer Service* on page xvii.

Displaying member service information

Click `Group`, expand `Members`, select the member name, and then click the `Service tab`.

In the Member Service window, component and disk information is specific to your array model.

You can display specific information about member hardware (for example, a component model, revision, or serial number). You can also display the Dell service tag identification for each individual array (not available on all arrays).

Collecting diagnostic information

In rare cases, an event might occur that only your PS Series support provider can correct. Your support provider might instruct you to collect encrypted diagnostic information from one or more group members. The group automatically sends this information to your support provider by using multiple e-mail messages.

Note: Do not collect diagnostic information unless instructed by your support provider.

1. Configure the group to use an SMTP server:

   a. Click `Group`, then `Group Configuration`, and then the `Notifications` tab.

   b. In the E-Mail Event Notifications panel, under E-Mail Configuration Settings, click `Add`, enter the IP address for an SMTP server (or e-mail relay), and click `OK`.

      Use the *`ip_address:port`* format to specify a port number other than the default (25).
      You can enter up to three IP addresses. The group uses one SMTP server or e-mail relay at any time. The first server you specify is the default server. The group uses the other servers in the order specified, if the default server is not available. Click the up and down arrows to change the order.

   c. Optionally, in the `Sender e-mail address` field, enter the address that appears in the message "From" field.

   d. Click `Save all changes` (Control+S) in the Group Notifications window.

2. Optionally, configure E-Mail Home notification. See *Configuring E-Mail home* on page 14-4. The support provider address for E-Mail Home receives the reports.

3. Click `Tools` and then `Diagnostic reports`.

4. In the Generate and E-Mail Diagnostics dialog box:

   - Select the member(s) for which you want to generate reports.

- Select whether to send reports to your support provider (requires E-Mail Home), one or two e-mail addresses (separated by a comma), or both.

Note: Do not change the default settings unless your support provider instructs you to change them.

5. Click OK.

To monitor progress while the group generates diagnostic reports, check the Alarms and Operations panel.

Appendix A Legal notices

This appendix lists the third-party copyrights for software used in the PS Series product.

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.

Glossary

This glossary defines the storage technology terminology that is specific to EqualLogic. If a term has unique meaning in the context of hardware or of a specific software application, that context is indicated.

See The SNIA Dictionary (http://www.snia.org/education/dictionary/) for definitions of any industry-standard storage terms used in this manual.

.bcd

(Auto-Snapshot Manager) The file extension used by ASM to identify a Smart Copy backup document.

.pvss

(Auto-Snapshot Manager) The file extension for a backup document relating to a mounted Smart Copy Set (post-VSS).

access control list (ACL)

A list of permissions attached to an object such as a storage volume. See access control record.

access control record

Means by which you restrict/control access to a PS Series volume. To ensure that only authorized computers and users can access a PS Series volume (iSCSI target), use access control records. You can restrict volume access by using any combination of specific IP addresses, iSCSI initiators, or CHAP user name and password (secret).

access credentials

Identity information that is checked against access control records. A computer must match the credentials specified in an access control record in order to connect to a volume. See access control record.

actions pane

(Auto-Snapshot Manager) The part of the ASM GUI main window that itemizes operations (actions). When you select a node in the Console tree, available operations for that node appear in the Actions Pane.

active control module

(hardware) In a dual control module array, the control module that is actively serving I/O on a network. If it ceases to function, it fails over to the secondary control module. See secondary control module.

application component

(Auto-Snapshot Manager) Any part of an application that ASM supports for Smart Copy operations.

application component (node)

(Auto-Snapshot Manager) Objects in the ASM console tree that represent components of software applications, such as SQL Server or Microsoft Exchange.

applications

(Auto-Snapshot Manager) Installed applications for which a VSS writer is available, such as Microsoft Exchange or SQL Server.

applications master node

(Auto-Snapshot Manager) The location of applications and their components in the ASM console tree.

apply logs

(Auto-Snapshot Manager) An ASM option that enables you to apply database transaction logs manually to a restored database.

array

(hardware) A PS Series storage array is a completely self-contained storage unit that includes multiple disk drives configured in a highly available RAID set with dual control modules and redundant power supplies and cooling modules.

array serial number

(Auto-Snapshot Manager) A unique PS Series array identification string that is encoded in the array's hardware. See service tag.

ASM/ME

(Auto-Snapshot Manager) Auto-Snapshot Manager/Microsoft Edition, a snap-in console application for the Microsoft Management Console that enables you to administer Smart Copies.

ASM/VE

(Auto-Snapshot Manager) Auto-Snapshot Manager/VMware Edition, a web-based application that works with VMware virtual environments to enable you to administer Smart Copies.

ASMCLI

(Auto-Snapshot Manager) A set of Smart Copy operation commands that you can execute at the Windows command prompt.

automatic RAID configuration

Internal process that configures the user-selected RAID policy on a member's disk array.

backup document

(Auto-Snapshot Manager) A file residing on the host that describes a Smart Copy on the PS Series array.

backup type

(Auto-Snapshot Manager) Specifies the backup behavior type that you want to create, either copy or full. Microsoft Exchange Storage Groups support only the copy backup type.

backup validation on startup

(Auto-Snapshot Manager) Automatically validate Smart Copies when ASM is started.

bandwidth

Rate at which an I/O subsystem or component can transfer bytes of data. Also called the transfer rate.

base volume

1. (Auto-Snapshot Manager) A PS series array volume mounted on the computer and reachable through its Windows-assigned drive letter (such as G:) or mount point.

2. (Group Manager) A volume that has snapshots. Snapshots depend on the base volume. If the base volume is destroyed, the snapshots have been removed.

broken smart copies

(Auto-Snapshot Manager) Smart Copies that ASM can no longer validate.

CHAP

Challenge Handshake Authentication Protocol, a network login protocol that uses an encrypted challenge-response mechanism. Used to limit access to volumes and snapshots to hosts that supply the correct account name and password. CHAP is also used for login/administrator accounts. See access credentials.

CHAP account

An account that uses CHAP configured on an external RADIUS server.

CHAP properties

(Auto-Snapshot Manager) An ASM configuration option that enables you to specify CHAP credentials for VSS or VDS access to groups and for computer access to Smart Copies for importing.

checksum verification

The process of verifying the integrity of Microsoft Exchange Smart Copies. You use the Windows eseutil.exe, a database maintenance utility.

cloning

The process of creating a new copy of an object such as a volume. The new object is the same type and has the same contents as the original. Contrast with a *thin clone*.

collection

(Auto-Snapshot Manager) Related groups of volumes or application components. These objects are represented by nodes in the ASM Console Tree under the Collections master node. Create collections of related volumes that you copy frequently. This ensures that ASM creates all the relevant Smart Copies simultaneously in one set. Create, modify, or delete a collection, create a Smart Copy Set for the collection, or configure a schedule for the collection.

command generation

(Auto-Snapshot Manager) The process of using the ASM GUI to generate an ASMCLI command.

console pane

(Auto-Snapshot Manager) The section of the ASM GUI that contains the console tree. This pane contains a collapsing, branched structure of related groups of clickable objects on which you can perform many different operations. See console tree.

console tree

(Auto-Snapshot Manager) A hierarchical structure of branched nodes representing objects on which ASM can perform operations. Nodes represent objects such as applications, volumes, and collections. The ASM console tree consists of related groups of objects (nodes) organized in a branching tree structure. Depending on the status of a node, you are presented with a menu of actions in the Actions Pane. See console pane.

control module

(hardware) The processor and interface component in a PS Series array. A control module contains the PS Series firmware in flash memory and provides temporary power continuity for data stored in cache memory. It has multiple network interfaces and an optional serial port. An array can contain two hot-swappable, dual redundant controllers. The active control module serves I/O, while the secondary control module mirrors data in its cache.

cooling module

(hardware) Hot-swappable hardware component, optionally integrated with a power supply, that provides cooling to a PS Series array. Arrays are shipped with redundant cooling modules. An array can continue to operate if one cooling module fails.

custom snapshot collection

(Auto-Snapshot Manager) One or more snapshots created at the same time through a multi-volume snapshot operation.

defer verification

(Auto-Snapshot Manager) When creating a replica, this operation causes ASM to defer Checksum Verification and Soft Recovery to a later time. Invoke the procedure manually or create a schedule at some future time.

delegated space

(Group Manager) Space on a group set aside to store received replicas for a partner.

demote

(Group Manager) Convert a volume in a replication configuration into a replica set.

device-specific module (DSM)

(Host Integration Tools) A plug-in for Microsoft Windows device driver module. For multipath implementation on PS Series arrays, you use EqualLogic Multipath I/O DSM in conjunction with Microsoft MPIO. It provides the EHCMservice.exe user mode Windows service, and the eqldsm.sys kernel mode driver.

discovery

Requesting from the target portal a list of accessible iSCSI targets (for example, volumes or snapshots) making those targets available for use.

DSM

See device-specific module (DSM).

eqlvdshwpriv

(hardware) A VDS service that runs on Windows and is specific to EqualLogic PS Series Group operations.

eqlvss

(hardware) A VSS service that runs on Windows and is specific to EqualLogic PS Series Group operations.

failback

(Group Manager) Replicating only the volume changes (delta) from the secondary group to the primary group and then returning to original replication configuration.

failback baseline

Date and time at which the data in the failback snapshot is identical to the data represented by the most recent replica.

failback replica set

(Group Manager) Temporary replica set created by demoting a volume as part of a failback operation. You can also create failback thin clone replica sets. See demote.

failback snapshot (baseline)

(Group Manager) A snapshot on the primary group containing the same data as the most recent complete replica, defining the failback baseline. A failback snapshot enables you to fail back to the primary group by replicating only the changes made to the recovery volume.

fan tray

(hardware) See cooling module.

global Smart Copy access

(Auto-Snapshot Manager) Refers to access controls that allow other computers to access (import) Smart Copies created on a particular computer.

global verification task

(Auto-Snapshot Manager) A scheduled background activity that you can run from any designated user account. The Global Verification Task performs Checksum Verification and Soft Recovery processing on Exchange Smart Copies.

global verification window

(Auto-Snapshot Manager) A core time period in which ASM can perform Checksum Verification and Soft Recovery on Smart Copies of Microsoft Exchange Storage Groups and mailbox databases. You typically specify a range of time that corresponds with a period of low system usage (off-peak times) to make best use of server resources.

group

See PS Series group.

group access

1. (Auto-Snapshot Manager).The process of enabling computer access to a PS Series Group by configuring and supplying credentials. See access credentials and CHAP.
2. (Group Manager) Access to the Group Manager UIs for management purposes. Access the GUI or CLI though the network. You can access the CLI through the optional serial port on the controller.
3. (hardware) Access to the group storage. iSCSI initiators access group iSCSI targets through the group IP address (discovery address). Access to a specific target is controlled through the access controls assigned to the target.

group administrator

An account on a PS Series group that has permission to manage all features and objects in a PS Series group, including configuring replication partners. The default group administrator account is grpadmin. See pool administrator.

group IP address

The network address that iSCSI initiators use to discover iSCSI targets and administrators use to access the group. See management IP address.

group member

See member.

group name

A unique identifier assigned to a group.

Host Integration Tools

(HIT) A suite of applications that enable you to configure and manage an array. It includes ASM/ME, DSM (Multipath I/O Device Specific Module), and RSW (Remote Setup Wizard).

hot-swap

(hardware) Removing a redundant component and installing a replacement while the array is running.

imported Smart Copy credentials

(Auto-Snapshot Manager) When a computer imports a Smart Copy, it must automatically present default credentials that match one of the Smart Copy's access control records.

iSCSI host bus adapter (HBA)

(hardware) An iSCSI initiator implemented as a physical I/O adapter through which a computer connects to a physical storage device such as a volume.

iSCSI initiator

The hardware or software component in a computer that starts the transfer of information to or from an iSCSI target. See iSCSI host bus adapter (HBA).

iSCSI portal verification

(Auto-Snapshot Manager) Verifying that ASM can connect to arrays to which it previously connected when last active.

iSCSI target

An iSCSI block storage device that you access through an iSCSI initiator. The volumes and snapshots in a PS Series group appear on the network as individual iSCSI targets.

iSCSI target discovery

An iSCSI protocol exchange that finds the available iSCSI targets from a target portal or discovery address.

jumbo frames

Ethernet frames capable of more than 1,500 bytes of payload (MTU). Enabling jumbo frames might improve performance on certain configurations.

keep count

1. (Auto-Snapshot Manager) The maximum number of snapshots or replicas retained by a Smart Copy schedule.
2. (Group Manager) The user-established limit on the number of snapshots or replicas created by using a schedule on the PS Series group.

latency

The time required to complete a specific I/O operation.

load balancing

Automatic distribution of I/O across resources to improve performance.

local replication reserve

(Group Manager) Storage space on a primary group that is used to record changes to a volume when replication is configured, and optionally to store a failback snapshot for a volume.

management IP address

In a group with a management network configured and enabled, an address used exclusively to log into the Group Manager GUI or CLI.

management network

An optional management network separates iSCSI traffic (volume I/O) from management traffic (GUI and CLI sessions, and other group management communications and intergroup operations).

manual restore

(Auto-Snapshot Manager) Mounting a Smart Copy and manually restoring data items.

manual transfer replication

(Host Integration Tools, Group Manager) Replication done through transportable media instead of over a network. Used in cases where the network link between replication partners is too slow or otherwise unsuitable for transferring large amounts of data.

Manual Transfer Utility

A stand-alone utility from EqualLogic that performs volume replication using transportable media, instead of the network. The utility has both graphical and command line user interfaces.

member

A PS Series array configured into a PS Series group. Groups can have several members.

member name

The unique name used to identify a specific member within a group.

membership password

The password required to add an array to a group, making it a member of the group.

merging pools

(Group Manager) The process of moving all the members and volumes from a source pool to a destination pool, deleting the source pool on completion.

mount

(Auto-Snapshot Manager) To create a connection to an iSCSI volume (clone, replica, or snapshot) and make its file system accessible to the operating environment.

MPIO

Acronym for multipath I/O. Multiple connections from an iSCSI initiator to targets on a PS Series Group over the network to provide redundancy and enhance performance. See device-specific module (DSM).

MPIO properties tab

(HIT) A EqualLogic-specific tab on the iSCSI Initiator properties page that provides status information about multipathing sessions. The ehcmn.log file is a rotating log file (such as ehcm0.log) containing the data displayed in the MPIO properties tab. See MPIO.

notification

(Auto-Snapshot Manager, Group Manager) The method that a group uses to inform you of significant events through e-mail, remote syslog files, and SNMP traps.

path failover

(hardware, Host Integration Tools) Relocating data traffic from a failed network path to a functional network path. This can occur automatically if the computer's software and hardware is configured for failover. MPIO provides server-side path failover.

path uptime

The elapsed time during which a session is active, displayed in the MPIO properties tab.

pool

Storage space provided by one to four group members. You assign volumes to a specific pool and load balancing operates only within pools. See load balancing and merging pools

pool administrator

(Group Manager) An account on a PS Series group that has permission to manage objects only in a specific pool or set of pools for a group. Compare to group administrator.

power supply

(hardware) Hot-swappable hardware component, sometimes integrated with a cooling module, that enables you to connect a PS Series array to a source of power. Arrays are shipped with redundant power supplies. An array can continue to operate if one power supply fails. Dell recommends that you connect power supplies to different sources of power, preferably on separate circuits.

primary group

(Group Manager) In a replication partnership the group containing the original volume. See secondary group.

primary volume

(Group Manager) A volume configured for replication to a replication partner.

promote

(Group Manager) To convert a replica set in a replication configuration into a volume. See demote.

PS Series array

A single EqualLogic iSCSI storage unit, usually configured as a PS Series Group. You can join multiple PS Series arrays into a larger PS Series Group and manage them as a single iSCSI SAN.

PS Series group

An iSCSI storage entity comprised of one or more PS Series storage arrays that you access through a single IP address and manage as a storage area network (SAN).

Queue-depth reporting

(SAN HQ) The average number of outstanding I/O operations at the start of each incoming I/O operation.

RAID policy

The type of RAID level (such as RAID 10 or RAID 6) that you configure for a member, coupled with the sparing policy (spares or no spares).

read-only account

(Group Manager) An administration account that only provides read-only access to group information.

recovery volume

Temporary volume created by promoting an inbound replica set as part of a failover operation. You can also create recovery template volumes and recovery thin clones. See promote.

replica

A point-in-time representation of a PS Series volume. The original volume and its replica are located on different PS Series groups (replication partners) potentially separated at some geographical distance to facilitate disaster tolerance.

replica collection

(Group Manager) The set of replicas resulting from each replication of a volume collection.

replica collection set

(Group Manager) The set of replica collections for a volume collection.

replication partner

(Group Manager) A group that is configured to send or receive replicas from another partner.

replica reserve

(Group Manager) Portion of the delegated space on a replication partner that is set aside for the replica sets for a specific volume. You configure the replica reserve for the volume on the primary group, but the actual replica reserve is on the secondary group.

replica set

(Group Manager) Set of complete replicas for a volume, template volume, or thin clone volume.

replication

(Group Manager) Copying volume data (only deltas) from the primary group, where the volume is stored, to the secondary group. Groups can be an unlimited distance apart. You can recover data from the secondary group, if necessary.

restore

1. (Auto-Snapshot Manager) The process of recovering data from a Smart Copy.

2. (Group Manager) The process of restoring the contents of a volume from a snapshot.

RSW

(HIT) Remote Setup Wizard, a graphical user interface (GUI) that enables you to configure a PS-Series array after you install the Host Integration Tools.

SAN HeadQuarters

(SAN HQ) Enables you to monitor multiple PS Series groups from a single graphical interface. It gathers and formats performance data and other important group information.

secondary control module

(hardware) Mirrors cache data from the active control module. If the active control module ceases to function, the secondary takes over network operations. See active control module.

secondary group

(Group Manager) In a replication configuration, the group that receives replicas of a source volume. See primary group.

service tag

(Group Manager) A unique ID assigned by Dell to particular equipment, for use by customer service.

sessions

(Group Manager) In an multipath configuration, the number of connections made to targets that span multiple members. The number of sessions can be configured to manage bandwith use.

shrink (volume)

(Group Manager) Decreasing the reported size of a volume.

Smart Copy

(Auto-Snapshot Manager) Point-in-time, application-consistent copy of objects in a PS Series group. Smart Copies can be of type snapshot, clone, or replica, depending on the edition of Auto-Snapshot Manager that you are using.

snapshot

A point-in-time representation of a PS Series iSCSI volume. Seen on the network as an iSCSI target. This is maintained in an array as deltas from the original volume.

snapshot collection

A set of snapshots resulting from a snapshot operation on a volume collection. See volume collection.

spare disk

(hardware) An unused disk in a PS Series array that is used automatically to replace a failed disk.

storage pool

(Group Manager) See pool.

template volume

Read-only volume from which you create thin clones.

thin clone

Volume that shares space with a template volume. Thin clones provide an efficient use of storage space for configurations with multiple volumes that have a large amount of common data.

thin provisioning

(Group Manager) The process optimizing use of storage space in a group through over-allocation. An object (such as a volume) is attributed less physical space than is reported by the group to any computer that is connected to the volume.

torn Smart Copy

(Auto-Snapshot Manager) The resulting Smart Copy that contains only partial data. The partial data set is referred to as torn because it does not contain all the files in a particular data set. This situation can occur when you attempt to do a Smart Copy of data sets that span multiple volumes.

transportable

(Auto-Snapshot Manager) A characteristic of a Smart Copy that enables it to be created on one computer and then moved to another computer to be used.

unmanaged space

Delegated space capacity on the secondary group that is no longer accessible from the primary group.

vacate

(Group Manager) To remove a member from a group while the member remains online.

VDS provider

(hardware) The EqualLogic VDS provider is a component of the Host Integration Tools that enables you use Microsoft Storage Manager for SANS to create and manage volumes in a PS Series Group.

volume

Storage allocated by a PS Series group as an addressable iSCSI target.

volume collection

(Group Manager) A number of volumes grouped together for purposes of performing operations on the volumes simultaneously. See also snapshot collection and replica collection.

volume administrator

An account on a PS Series group that has permission to manage a quota of storage in one or more pools. A volume administrator can perform any volume operations, including replication, within their quota limit.

volume reserve

(Group Manager) Amount of space allocated to a volume from free pool space. Without thin provisioning, volume reserve is the same as reported size.

vss-control volume

(Auto-Snapshot Manager) A special logical volume that enables Microsoft VSS/VDS services to communicate with a PS Series array. The vss-control volume appears as a non-initialized disk in the Windows Disk Management interface.

XMLlogsize

(hardware) The variable that controls the size of ehcmn.log.

Index

W