



PS SERIES STORAGE ARRAYS GROUP ADMINISTRATION

PS Series Firmware Version 3.3



Copyright 2008 EqualLogic, Inc. All rights reserved.

EqualLogic is a registered trademark.

Dell is a trademark of Dell, Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Information in this document is subject to change without notice.

Reproduction in any manner whatsoever without the written permission is strictly forbidden.

February 2008

Part Number: 110-5016-R1

Table of Contents

Preface	xi
Audience.....	xi
Organization	xi
Conventions.....	xii
Documentation	xii
Technical Support and Customer Service	xiii
 1 Introduction.....	 1-1
Storage Solutions for All Enterprises.....	1-1
PS Series Groups	1-2
High-End Features in an Affordable iSCSI SAN.....	1-3
Modular Hardware.....	1-3
Seamless Online Scalability	1-4
Interoperability	1-5
Easy Setup and Management.....	1-5
Automatic SAN Operation	1-5
Robust Security for Group Administration	1-6
Robust Security for Data Access	1-7
Advanced Functionality at No Extra Cost.....	1-7
Host Integration Tools.....	1-8
 2 Getting Started	 2-1
Essential Setup Steps.....	2-1
Configuring an Array and Creating a Group.....	2-2
Adding Members to an Existing Group	2-5
Using the Group Manager Graphical User Interface	2-7
Controlling Group Manager GUI Operation.....	2-11
Securing GUI Access.....	2-11
Setting GUI User Preferences.....	2-11
General Policies.....	2-12
Communication Policies	2-13
Alarm Policies	2-14
Data Validation and Debugging Policies	2-15
Obtaining Online Help.....	2-16
Running the GUI on a Local System.....	2-17
Using the Group Manager Command Line Interface.....	2-19
Where to Go Next	2-20

3 Group Customization	3-1
Displaying the Group Configuration and Status	3-1
Modifying the Group Network Configuration	3-5
Modifying the Date, Time, and Time Zone	3-6
Controlling GUI and CLI Management Access	3-7
Managing Group Administration Accounts	3-9
Using Local Administration Accounts	3-9
Using External RADIUS Accounts	3-13
Setting Up External RADIUS Authentication Servers	3-13
Configuring RADIUS Authentication Servers	3-16
Configuring RADIUS Accounting Servers	3-18
Setting Up Event Notification	3-20
Configuring E-Mail Event Notification	3-21
Configuring E-Mail Home Notification	3-23
Configuring Syslog Event Notification	3-24
Displaying Live Informational Messages	3-24
Using CHAP to Control iSCSI Access to Volumes	3-25
Using Local CHAP Accounts for Initiator Authentication	3-26
Using External RADIUS Servers for Initiator Authentication	3-28
Configuring Target Authentication	3-29
Using iSNS for iSCSI Device Management	3-30
Using SNMP to Monitor a Group	3-31
Configuring SNMP Read-Only Access	3-31
Sending SNMP Traps to a Destination	3-32
Modifying Access Controls for Windows Services	3-35
Modifying the Group-Wide Volume Settings	3-36
Controlling Performance Load Balancing	3-38
Preventing Member Firmware Downgrades	3-40
 4 Storage Pool Management	 4-1
Using Storage Pools to Separate SAN Storage	4-1
Planning Storage Pools	4-4
Creating a Storage Pool	4-5
Displaying Storage Pool Configuration and Status	4-6
Modifying a Storage Pool Name or Description	4-10
Merging Storage Pools	4-10
Deleting a Storage Pool	4-11

5 Volume Management.....	5-1
Introduction to Volumes.....	5-1
Creating Volumes.....	5-3
Understanding Thin Provisioning	5-8
Displaying Volume Status and Configuration	5-11
Managing Volume Collections.....	5-15
Creating a Volume Collection	5-15
Displaying Volume Collection Status	5-17
Modifying a Volume Collection.....	5-22
Deleting a Volume Collection	5-23
Managing Access Controls for Volumes and Snapshots	5-23
Creating Access Control Records.....	5-24
Displaying Access Control Records	5-26
Modifying an Access Control Record	5-26
Deleting an Access Control Record.....	5-27
Connecting Hosts to Volumes and Snapshots.....	5-27
Modifying Volume Settings.....	5-29
Moving a Volume Across Pools.....	5-29
Increasing the Size of a Volume.....	5-31
Decreasing the Size of a Volume	5-33
Modifying a Volume Name, Description, and Alias	5-33
Enabling and Disabling Thin Provisioning	5-34
Setting a Volume RAID Preference	5-35
Binding a Volume to a Specific Member	5-37
Cloning Volumes.....	5-38
Managing Schedules for Snapshots and Replicas	5-41
Setting a Volume Offline or Online	5-48
Setting Volume Permission.....	5-48
Deleting Volumes.....	5-49
Handling Lost Blocks.....	5-49
 6 Snapshot Management	 6-1
Introduction to Snapshots.....	6-1
Modifying Snapshot Settings	6-3
Creating Snapshots.....	6-5
Displaying Snapshot Configuration and Status.....	6-6
Setting a Snapshot Online or Offline	6-11
Setting Snapshot Permission	6-12
Modifying a Snapshot Name, Description, and Alias.....	6-12

Scheduling Snapshots.....	6-14
Restoring a Volume from a Snapshot	6-14
Cloning a Snapshot.....	6-15
Deleting a Snapshot.....	6-16
Managing Snapshot Collections.....	6-16
Creating a Snapshot Collection	6-17
Creating a Custom Snapshot Collection	6-18
Displaying Snapshot Collection Status.....	6-20
Modifying a Snapshot Collection.....	6-23
Deleting a Snapshot Collection	6-25
7 Basic Replication	7-1
Introduction to Replication.....	7-1
Replication Configurations.....	7-3
Understanding Replication Space Requirements	7-5
Sizing Delegated Space and the Replica Reserve	7-5
Sizing the Local Replication Reserve	7-7
Steps for Replicating Volumes.....	7-9
Managing Replication Partners	7-11
Replication Partner Requirements	7-11
Configuring Replication Partners	7-11
Modifying a Replication Partner	7-17
Deleting a Replication Partner.....	7-18
Configuring a Volume or Volume Collection for Replication.....	7-19
Disabling Replication on a Volume or Volume Collection	7-24
Creating Replicas	7-24
Scheduling Replication	7-25
Displaying Replication Status	7-26
Understanding Replication Status	7-26
Displaying the Replication Configuration for a Volume.....	7-29
Displaying Replica Collections	7-31
Displaying Replication Activity Between Partners	7-32
Displaying Outbound Replication Status	7-35
Displaying Inbound Replication Status	7-38
Pausing and Resuming Replication.....	7-41
Cancelling an In-Progress Replication.....	7-41
Cloning an Inbound Replica	7-42
Deleting Replicas and Replica Sets.....	7-45
Deleting Outbound Replicas.....	7-45

Deleting Inbound Replicas	7-46
Deleting a Replica Collection	7-47
8 Advanced Replication	8-1
Performing Manual Transfer Replication	8-1
Installing the Manual Transfer Utility	8-2
Using the Manual Transfer Utility GUI	8-3
Creating Transfer Files Using the GUI	8-4
Loading Transfer Files Using the GUI	8-11
Managing Completed Operations Using the GUI	8-17
Using the Manual Transfer Utility CLI	8-19
Creating Transfer Files Using the CLI	8-20
Creating Additional Transfer Files Using the CLI	8-24
Loading Transfer Files Using the CLI	8-25
Recovering Data on the Secondary Group	8-28
Procedures for Recovering Data	8-29
Steps for Hosting a Volume for Backup Purposes	8-30
Steps for Failback	8-31
Steps for Switching Roles in a Replication Configuration	8-34
Temporarily Promoting an Inbound Replica Set to a Recovery Volume	8-36
Demoting a Recovery Volume to an Inbound Replica Set	8-40
Making an Inbound Replica Set Promotion Permanent	8-42
Permanently Promoting an Inbound Replica Set to a Volume	8-44
Demoting a Volume to a Failback Replica Set	8-47
Promoting a Failback Replica Set to a Volume	8-48
Changing the Pool for a Failback Replica Set	8-50
Making a Volume Demotion Permanent	8-51
9 Member Management	9-1
Introduction to Member RAID Policies	9-1
Setting the RAID Policy and Pool for a New Member	9-3
Displaying Member Configuration and Status	9-6
Managing the Disk Array	9-12
Displaying the Disk Configuration	9-12
Changing the Member RAID Policy	9-14
Installing and Removing Disks	9-16
Handling Failed Disks	9-17
Managing the Member Network Configuration	9-18
Configuring a Network Interface	9-19

Displaying Network Interface Status	9-20
Enabling or Disabling a Network Interface	9-21
Deconfiguring a Network Interface	9-21
Modifying the Default Gateway	9-22
Managing Control Modules	9-23
Displaying Control Module Status	9-23
Setting Control Module Cache Modes	9-25
Handling Failed Control Modules	9-27
Moving a Member Across Pools	9-27
Modifying a Member Name and Description	9-29
Modifying the Group Membership Password	9-30
Removing a Member from a Group	9-30
Using LEDs to Identify a Member	9-31
Displaying Member Enclosure Status	9-32
Displaying Member Connections	9-35
Obtaining Member Service Information	9-36
10 Diagnosing Problems	10-1
Monitoring Events	10-1
Displaying Events	10-1
Accessing the Event Log File on a Remote System	10-4
Accessing Events Sent to an E-Mail Address	10-4
Examining SNMP Traps	10-4
Monitoring Alarms and Action Items	10-4
Monitoring the Group	10-7
Using the Group Monitoring Windows	10-7
Displaying iSCSI Connections	10-7
Displaying Schedules	10-8
Displaying Replication Information	10-9
Displaying Login Information	10-10
Displaying Member and Volume Moving Operations	10-11
Monitoring Space	10-12
Monitoring Replication	10-13
Monitoring Volumes and Snapshots	10-14
Monitoring Hardware	10-14
Monitoring Performance Statistics	10-15
Using MRTG to Monitor I/O and Performance	10-18

A Glossary	A-1
B Legal Notices	B-1
Index	Index-1

Preface

This manual describes how to use the Group Manager graphical user interface (GUI) to manage a PS Series group, a self-managing, iSCSI storage area network (SAN) that is affordable and easy to use, regardless of scale. The information in this document is useful to individuals who want to understand overall group features and operations.

For information about using the command line interface (CLI) to manage a group and individual PS Series storage arrays, see the *CLI Reference* manual.

Audience

This manual is designed for administrators responsible for managing a PS Series group and also for individuals who want to learn more about group operation. Administrators are not required to have extensive network or storage system experience. However, it is useful to understand:

- Basic networking concepts
- Current network environment
- User disk storage requirements
- RAID configurations
- Disk storage management

Note: Although examples are provided showing the use of PS Series arrays in some common network configurations, detailed information about setting up a network is beyond the scope of this material.

Organization

The organization is as follows:

- Chapter 1, *Introduction* provides an introduction to arrays, group features, and functionality.
- Chapter 2, *Getting Started* describes the steps that allow you to quickly set up and start using a group, including creating a group and adding members. Also includes a brief discussion of post-setup tasks and an introduction to the management interfaces.
- Chapter 3, *Group Customization* describes how to modify the group configuration.

- Chapter 4, *Storage Pool Management* describes how to separate storage in a multi-member group.
- Chapter 5, *Volume Management* describes managing volumes and volume collections.
- Chapter 6, *Snapshot Management* describes how to manage volume and volume collection snapshots.
- Chapter 7, *Basic Replication* describes how to perform basic volume replication tasks.
- Chapter 8, *Advanced Replication* describes how to perform advanced replication tasks, including manual transfer replication and recovering data.
- Chapter 9, *Member Management* describes how to manage group members.
- Chapter 10, *Diagnosing Problems* describes how to monitor events, interpret status, diagnose problems, and monitor performance statistics.
- Appendix A, *Glossary* lists and defines glossary terms used in the documentation.
- Appendix B, *Legal Notices* lists copyright and trademark information.

Conventions

Documentation conventions are shown in the following table.

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 desired object.
Volumes > Access tab	When navigating the GUI, indicates that the user should click on <i>Volumes</i> in the far left column and then click the <i>Access</i> tab in the window that appears.

Documentation

For detailed information about PS Series arrays, groups, and volumes, see the following documentation:

- *Release Notes*. Provides the latest information about PS Series arrays.
- *QuickStart*. Describes how to set up the array hardware and create a PS Series group. Be sure to use the manual for your array model.
- *Group Administration*. Describes how to use the Group Manager graphical user interface (GUI) to manage a PS Series group. This manual provides comprehensive information about product concepts and procedures.

- *CLI Reference*. Describes how to use the Group Manager command line interface (CLI) to manage a PS Series group and individual arrays.
- *Hardware Maintenance*. Provides information about maintaining the array hardware. Be sure to use the manual for your array model.
- Online help. In the Group Manager GUI, expand `Tools` in the far left panel and then click `Online Help` for help on both the GUI and the CLI.

The *QuickStart* and *Hardware Maintenance* manuals and the *Release Notes* are printed and shipped with the product. They are also located on the PS Series documentation CD-ROM that is shipped with the product, along with the *Group Administration* and *CLI Reference* manuals and the Group Manager online help.

Documentation for the Host Integration Tools for Microsoft® Windows® systems is located on the Host Integration Tools CD-ROM that is shipped with the product.

Localized documentation is available on the localization CD-ROM that is shipped with the product in certain countries.

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 may not be available in your area.

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

Note: If you do not have an Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog.

To contact Dell for sales, technical support, or customer service issues:

1. Visit `support.dell.com`.
2. Verify your country or region in the Choose A Country/Region drop-down menu at the bottom of the window.
3. Click `Contact Us` on the left side of the window.

4. Select the appropriate service or support link based on your need.
5. Choose the method of contacting Dell that is convenient for you.

Online Services

You can learn about Dell products and services on the following websites:

- www.dell.com/
- www.dell.com/ap/ (Asian/Pacific countries only)
- www.dell.com/jp (Japan only)
- www.euro.dell.com (Europe only)
- www.dell.com/la (Latin American countries)
- www.dell.ca (Canada only)

You can access Dell Support through the following websites:

- support.dell.com
- support.dell.com/EqualLogic
- support.jp.dell.com (Japan only)
- support.euro.dell.com (Europe only)

1 Introduction

PS Series storage arrays deliver the benefits of consolidated storage in a self-managing, IP storage area network (SAN) that is easy to use and affordable—regardless of its size. Featuring automated management and fast, flexible scalability, PS Series arrays can greatly reduce the total cost of storage acquisition and management.

Storage Solutions for All Enterprises

Today, large and small businesses alike are under pressure to manage rapidly-growing storage environments. Large amounts of data are still directly attached to servers—a configuration that is difficult to scale and results in lower service levels.

To use storage resources efficiently and economically, the best solution is to consolidate disparate storage configurations. By reducing points of management, you can make storage administration easier and improve service levels, while facilitating data provisioning, backups, and disaster recovery. The challenge has been to find a consolidated storage solution that is easy to deploy, as well as to operate and maintain.

With PS Series arrays, the benefits of storage consolidation are available to all enterprises. Grouping together one or more PS Series arrays enables you to manage them as a single storage system. Now, businesses can leverage their existing IP network infrastructure to deploy an easy-to-manage and full-featured iSCSI SAN.

To make storage management simple and efficient, PS Series arrays use virtualization technology that masks the underlying complexity of the storage configuration—saving administrators time and effort. An IP network is used not only to connect hosts and applications to storage, but also to connect arrays to each other, providing a communication mechanism through which arrays share configuration data and collaborate during automatic data provisioning and load balancing.

Self-managing PS Series arrays enable administrators to quickly and easily make configuration changes—without disrupting running applications. You can add more arrays to a group to seamlessly scale capacity and performance or replicate data over long distances for a simple, yet robust disaster recovery plan. With support for a variety of operating systems, PS Series arrays ensure that you can consolidate all your storage in a easy-to-grow and easy-to-use SAN.

PS Series Groups

A **PS Series group** is an iSCSI SAN comprised of one or more PS Series storage arrays (**members**) connected to an IP network and managed as a single system.

Each PS Series array has fully redundant hardware and up to three active network connections for maximum bandwidth. Integrated virtualization firmware provides:

- Seamless scalability
- Automatic RAID configuration and spare disk configuration
- Automatic network, performance, and capacity load balancing

In a single-member group, data is automatically spread across the array disks. Add more arrays to the group to expand SAN capacity and increase performance—with no disruption to users. Management overhead stays the same, regardless of SAN size.

After an array is added to a group and you choose a RAID level for the new member, the disks are automatically configured, and the space is immediately available for use. Data is spread across group members and automatically load-balanced as needed, based on capacity and performance metrics.

To access the storage space in a group, use the Group Manager graphical or command line user interface to create **volumes**, each of which has a size, access controls, and other attributes. A volume can be spread across multiple disks and is seen on the network as an **iSCSI target**.


Optionally, you can use **thin provisioning** to improve storage resource utilization by logically over-allocating SAN volumes. Thin provisioning enables you to satisfy current and future storage needs and eliminates the difficulty of adding storage to operating systems or applications that do not support online expansion.

A host needs only an industry-standard **iSCSI initiator** to access a volume in a group. Most major operating systems support an iSCSI initiator (in the form of an iSCSI driver or a host bus adapter), with a range of price and performance options. Once connected, a volume appears as a regular disk that can be formatted with the operating system's normal utilities. Even as members are added and removed from the group, volumes continue to be accessible through the same iSCSI targets; no host modifications are necessary.

In a group, you can place all the members and volumes in a single **storage pool**. Or, if you have multiple members, you can create additional pools and assign members and volumes to different pools. Pools allow you to organize storage according to usage, providing more control over resource allocation, while maintaining a single system management view.

You retain the advantages of storage consolidation while gaining the ability to easily segregate workloads.

Figure 1-1 shows a three-member group with multiple storage pools.



For backup purposes, you can create point-in-time **snapshots** of volumes. You can also **replicate** volumes across multiple groups for a simple, yet robust disaster recovery plan. In addition, **failback** functionality enables you to temporarily host a volume from a replication partner and then fail back to the original replication configuration.

You can also **clone** volumes, snapshots, and replicas to create new volumes.

High-End Features in an Affordable iSCSI SAN

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

Modular Hardware

A PS Series group can easily grow or shrink to accommodate workload requirements, so administrators purchase only the hardware they need when they need it. New product versions will fully interoperate with prior-generation arrays, protecting your initial investment.

Each PS Series array includes:

- **No-single-point-of-failure hardware.** Redundant, hot-swappable hardware components—disks, control modules, fans, and power supplies—provide the highest availability. Component failover and disk sparing occur automatically without user intervention or disrupting data availability. In addition, data in each array is protected with your choice of RAID technology.
- **High-performance control modules.** Each control module has three Gigabit Ethernet interfaces. You can configure multiple network interfaces for automatic failover. With dual control modules, data is mirrored across battery-backed write-back caches. If one control module fails, the other control module takes over—automatically and with no disruption to users.
- **Support for standard Ethernet networks.** Only one IP network connection is needed for array operation. You can configure up to three network interfaces for maximum bandwidth. There is no need to train administrators in unfamiliar and complex technologies like Fibre Channel. Also, costs are controlled due to the high volumes and intense vendor competition among Ethernet hardware vendors.
- **Easy and online expandability.** New arrays are guaranteed to work with older models, ensuring the viability of your initial purchase. You can expand SAN capacity with no disruption to users and application.
- **Optional management connectivity through a serial port.** A serial connection enables you to connect the array to a console or computer running a terminal emulator. This ensures that you can access a group if network connectivity is lost.
- **High capacity SAS or SATA disk drives.** Depending on your array control module type, you can install high-capacity disks to increase your storage potential.

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 automatically, and the new disk and network resources are immediately available for use. The additional control modules and caches also increase processing power. Meanwhile, volumes remain available with no impact on hosts and applications, and management overhead stays the same.

Interoperability

PS Series arrays are ideal for multi-platform, heterogeneous environments, eliminating the need to have a different storage system for each operating system or application. Because a PS Series group provides block-level storage, hosts need only a standards-compliant iSCSI initiator—available from most major operating systems—to access data. Once a host connects to a volume, it appears as a regular disk that can be formatted using the normal operating system utilities.

Easy Setup and Management

A simple setup utility lets you 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 using a serial connection between an array and a console or computer running a terminal emulator. Alternately, Windows systems can run the Remote Setup Wizard to configure an array and set up system access to the SAN. See *Host Integration Tools* on page 1-8 for more information.

Password-protected management accounts enable easy, secure access to the group. Graphical and command line user interfaces provide a single-system view of the storage. Administrators do not need multiple consoles to manage the storage. Using the Group Manager graphical user interface (GUI), creating and managing volumes, configuring security, and setting up event notification are point-and-click operations. The equivalent command line interface (CLI) can be accessed through telnet, SSH, or a serial connection and also used in scripts that are run on servers.

Automatic SAN Operation

In contrast to traditional storage management environments that require you to perform error-prone, manual tasks, a PS Series group handles complex tasks intelligently and without user intervention:

- **Automatic RAID configuration and data provisioning.** There is no need for administrators to manually create RAID sets or map data onto disks or individual control modules. Whether you are expanding a group or creating, expanding, or deleting volumes, the group handles storage allocation and load balancing within and across the arrays in the group. Host iSCSI access to volumes continues without interruption.

- **Automatic spare disk configuration and use.** A PS Series array can include spare disks, which are automatically configured and used in the event of a disk failure, enabling “hands-off” storage management.
- **Dynamic load balancing.** A PS Series group delivers three types of load balancing within the arrays in a storage pool. This allows the group to quickly detect and eliminate “hot spots” as the workload changes, with no user intervention or host disruption:
 - Capacity load balancing ensures that data is optimally distributed across spindles, caches, and controllers.
 - Network connections are load-balanced within and across members, minimizing network “hops” and maximizing bandwidth.
 - Performance load balancing enables volume data to be located where the RAID configuration is optimal for the volume, based on internal group performance metrics. For example, if a volume would perform best in a RAID 10 configuration, the group will attempt to place volume data on arrays that are configured with RAID 10.
- **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 files, or SNMP traps.

Robust Security for Group Administration

You can access the group for management purposes by using telnet, SSH, or standard or encrypted Web access. Separate password-protected administration accounts can be created, as needed, with either read-write or read-only privileges. You can use accounts configured and authenticated on the group or on a RADIUS (Remote Authorization Dial-in User Service) server.

You can create three types of administration accounts:

- Group administration accounts have full access to all members and volumes in the group. Only these accounts can add and remove members, create pools, and configure replication partners.
- Pool administration accounts can perform operations only within specific pools and optionally can view group information, but cannot make group-wide changes. Pool administrators also have access to the group monitoring information and event logs.
- Read-only accounts can only monitor group members and volumes.

Robust Security for Data Access

Host iSCSI access to volumes can be restricted according to IP address, iSCSI initiator name, or CHAP user name (configured and authenticated locally or on a RADIUS server). There is no need to understand complicated security technologies, such as Fibre Channel Switch Zoning or LUN Masking.

Advanced Functionality at No Extra Cost

PS Series arrays match ease-of-use with ease-of-doing-business. Advanced features are included in the standard array offering, so there is no need to purchase expensive add-ons. New features can be incorporated into an existing group through firmware updates.

For example, a PS Series group delivers:

- **Online scalability.** You can expand array capacity or overall group capacity while online and with no impact on users or availability. Hot-swappable hardware means that failed components can be replaced while data remains online.
- **Thin provisioning.** In the right environment, administrators can improve storage resource utilization by using thin provisioning to logically over-allocate SAN volumes. Thin provisioning enables you to meet current and future storage needs and eliminates the difficulty of adding storage to operating systems or applications that do not support online expansion.

To an application or operating system, a thin-provisioned volume appears to be fully allocated. However, the group initially allocates only a portion of the volume size. As the volume is used, more space is automatically allocated, with no disruption in availability. Increased usage triggers automatic notifications, so administrators can make adjustments as needed.

- **Cloning.** Cloning a volume creates a new volume with the same size and contents as the original volume. Cloning is commonly used in multiple server deployments. For example, a master image of a system can be created and then cloned for each server. Cloning can dramatically reduce overhead when deploying replicated servers, such as blade servers or Web servers.
- **Snapshots.** A snapshot is a point-in-time copy of volume data that protects against mistakes, viruses, or database corruption. Snapshots greatly simplify and improve the performance of backup and recovery operations. Both on-demand and scheduled snapshots are supported.

To recover data from a snapshot, you can set the snapshot online, restore the volume from a snapshot, or clone the snapshot. Cloning a snapshot creates a volume that contains the volume data that existed at the time the snapshot was created.

- **Replication.** A replica represents the contents of a volume at a specific point in time. By replicating volumes from one group to another, you can set up a simple, yet robust disaster recovery plan that protects against catastrophic events. Groups can be located in the same building or separated by an unlimited distance. Both on-demand and scheduled replicas are supported.

To recover data from a replica, you can clone a replica to create a new volume; replication will continue as usual. If a disaster occurs on the primary group, you can promote a replica set to a new volume and snapshots—temporarily (with the option of failing back to the primary group) or permanently.

- **Failback.** Failback functionality enables you to temporarily host the volume on the secondary group and later return to the original replication configuration. In some cases, you can fail back by replicating only the volume changes that occurred while the volume was hosted on the secondary group.
- **Scheduled operations.** As needed, you can set up schedules to create volume snapshots or replicas at a specific time in the future or on a regular basis.
- **Volume collections.** You can organize volumes into collections. This enables you to perform a snapshot or replication operation on all the volumes in the collection simultaneously, ensuring that the resulting set of snapshots or replicas are consistent.
- **Multiple storage pools.** In a PS Series group, storage can be divided into multiple pools. This allows you to organize storage according to usage, providing more control over resource allocation, while maintaining a single system management view. You retain the advantages of storage consolidation while gaining the ability to easily segregate different workloads.

Host Integration Tools

You can use host-based solutions that expand PS Series group capabilities at no extra cost. The Host Integration Tools for Microsoft® Windows® are located on the CD-ROM shipped with the array and also on the Technical Support website. Always check the website to ensure that you have the latest version of the Host Integration Tools kit and documentation.

2 Getting Started

To get started with your PS Series array, see the following topics:

- *Essential Setup Steps* on page 2-1
- *Configuring an Array and Creating a Group* on page 2-2
- *Adding Members to an Existing Group* on page 2-5
- *Using the Group Manager Graphical User Interface* on page 2-7
- *Controlling Group Manager GUI Operation* on page 2-11
- *Using the Group Manager Command Line Interface* on page 2-19
- *Where to Go Next* on page 2-20

Essential Setup Steps

The following steps describe the process for getting started with a PS Series array.

Step 1: Set up the array hardware. Following all the safely precautions, install the array in a rack and connect the required cables. When you turn on power to the array, ensure there are no errors and no red LEDs.

See the PS Series *QuickStart* or SAN Setup poster for your model array for hardware installation information.

Step 2: Configure the array and create or expand a group. See *Configuring an Array and Creating a Group* on page 2-2 or *Adding Members to an Existing Group* on page 2-5.

Step 3: Choose the RAID policy for the member. To complete the configuration, use the Group Manager GUI or CLI to choose the RAID policy for the member. The storage in the member will be available once you set the RAID policy.

See *Using the Group Manager Graphical User Interface* on page 2-7, *Using the Group Manager Command Line Interface* on page 2-19, and *Introduction to Member RAID Policies* on page 9-1.

Step 4: Create a volume. Use the Group Manager GUI or CLI to create a volume. You must assign a name, size, and access controls to the volume.

See *Creating Volumes* on page 5-3.

Step 5: Connect to the volume from a host system. A volume is seen on the network as an iSCSI target. A host system needs an industry-standard iSCSI initiator to connect to the volume. Once a host is connected, the volume appears as a regular disk.

See *Connecting Hosts to Volumes and Snapshots* on page 5-27.

After performing the essential steps, you can customize the group and use the complete range of group features. See *Where to Go Next* on page 2-20.

Configuring an Array and Creating a Group

Use the `setup` utility to configure an array and create a PS Series group with the array as the initial group member. First, be sure the array hardware is set up and there are no errors or red LEDs. For hardware installation information, see the PS Series *QuickStart* or SAN Setup poster for your model array.

Note: If you have a Microsoft Windows system, you can use the Remote Setup Wizard to configure an array, create a new group or expand an existing group, and set up connectivity between the system and group. The wizard is part of the Host Integration Tools kit, which is on the CD-ROM shipped with each array, and the Technical Support website.

When configuring an array and creating a group, be sure to follow the network requirements and recommendations listed in Table 2-1. For the latest networking information, see the related Technical Report on the Technical Support website.

Table 2-1: Network Requirements and Recommendations

Requirement or Recommendation	Description
Connectivity to the group IP address	Each group member is <i>required</i> to have at least one configured network interface on the same subnet as the group IP address.
Switch connectivity	For a single-subnet group in which the members' configured network interfaces are connected to multiple switches, network connectivity between the switches is <i>required</i> .
Access to the group IP address subnet	For a multi-subnet group, it is <i>recommended</i> that the configured network interfaces be able to access the subnet on which the group IP address resides.

Follow these steps to configure an array and create a group:

1. Gather the information described in Table 2-2 and Table 2-3. The `setup` utility will prompt for this information. Your network administrator should be able to provide the appropriate IP addresses for the member and group configuration.

Table 2-2: Member Configuration

Prompt	Description
Member name	Unique, descriptive name (fewer than 64 alphanumeric characters, no spaces, and the first character must be a letter or number), used only to identify and administer the array.
Network interface name	Name of a network interface (<code>eth0</code> , <code>eth1</code> , or <code>eth2</code>) that is physically connected to a functioning port on a network switch (using a network cable). You can connect and configure additional network interfaces on the array after it is added to a group.
IP address	Network address for the network interface specified above. Each member must have at least one network interface on the same subnet as the group IP address.
Netmask	Combines with the previous IP address to identify the subnet on which the network interface specified above resides. The default is <code>255.255.255.0</code> .
Default gateway (optional)	Network address for the device that connects subnets and forwards network traffic beyond the local network. A default gateway is needed only if you want the network interface specified above to be able to communicate outside the local network (for example, for host access to volumes). The default gateway must be on the same subnet as the network interface specified above.

Table 2-3: Group Configuration

Prompt	Description
Group name	Name of the group for administrative purposes (up to 64 alphanumeric characters and hyphens). The first character must be a letter or number.
Group IP address	Network address for the group. The group IP address is used for administrative and host access to volumes.

Table 2-3: Group Configuration (Continued)

Prompt	Description
Password for managing group membership	Password required when adding members to the group. The password must have 3 to 16 alphanumeric characters and is case-sensitive.
New password for the default <code>grpadmin</code> account	Password that will override the factory-set password (<code>grpadmin</code>) for the default <code>grpadmin</code> account. The password must have 3 to 16 alphanumeric characters and is case-sensitive.

- Set up a serial connection to the array. Use the serial cable shipped with the array. Make the connection to Serial Port 0 on the active control module (the ACT LED will be green) and to a console or computer running a terminal emulator. The serial connection must have the following characteristics:
 - 9600 baud
 - One STOP bit
 - No parity
 - 8 data bits
 - No hardware flow control
- Log in to the `grpadmin` account. Press the `[Enter]` key and, when prompted, enter the `grpadmin` account name and factory-set password, which is also `grpadmin`. Note that passwords are *not* echoed on the console.

```
Login: grpadmin
Password:
```

```

      Welcome to Group Manager
      Copyright 2001 - 2007 EqualLogic, Inc.

```

```

It appears that the storage array has not been configured.
Would you like to configure the array now? (y/n) [n]

```

- Enter `y` to automatically invoke the `setup` utility. You can also run `setup` later at the console prompt (`>`). At each `setup` prompt, enter a question mark (?) for help. Press the `[Enter]` key to accept a default (shown in brackets).
- Specify the member and group configuration information at the prompts. This is the information described in Table 2-2 and Table 2-3.

Once the array is successfully added to the new group, use the Group Manager GUI or CLI to set the RAID policy for the member.

The storage in the member will be available once you set the RAID policy. The disks are automatically configured according to the designated RAID level, with the appropriate number of spare disks. Until the RAID configuration completes, performance will not be optimal, but the group is fully operational.

See *Using the Group Manager Graphical User Interface* on page 2-7 and *Using the Group Manager Command Line Interface* on page 2-19 for information on the starting the Group Manager user interfaces. See *Introduction to Member RAID Policies* on page 9-1 for information on setting the RAID policy.

Adding Members to an Existing Group

Initially, a PS Series group has only one member. Although this configuration is fully functional, adding more members expands the capacity of the group, increases network bandwidth, and improves overall performance. As the group scales, there is no disruption to I/O operations. See the PS Series *Release Notes* for the maximum members in a group.

In addition, a multi-member group enables you to divide group space into multiple storage pools. See *Storage Pool Management* on page 4-1 for more information.

You can use the `setup` utility to configure an array and add it to an existing group. First, be sure the array hardware is set up and there are no errors or red LEDs. For hardware installation information, see the PS Series *QuickStart* or SAN Setup poster for your model array.

Note: If you have a Microsoft Windows system, you can use the Remote Setup Wizard to configure an array, add it to an existing group, and set up connectivity between the system and the group.

The wizard is part of the Host Integration Tools kit, which is on the CD-ROM shipped with each array and the Technical Support website.

Follow these steps to use the `setup` utility to add a member to an existing group:

1. Gather the member configuration described in Table 2-2. See your network administrator to obtain an IP address for the new group member.

Also, you will need the name and IP address for the group that the array will join, in addition to the password for managing group membership.

Be sure to follow the network requirements and recommendations listed in Table 2-1. For the latest networking recommendations, see the related Technical Report on the Technical Support website.

2. Set up a serial connection to the array. Use the serial cable configuration that shipped with the array and connect the cable to Serial Port 0 on the active control module (ACT LED will be green).

The serial connection must have the following characteristics:

- 9600 baud
- One STOP bit
- No parity
- 8 data bits
- No hardware flow control

3. Log in to the `grpadmin` account. Press the `[Enter]` key and, when prompted, enter the `grpadmin` account name and factory-set password, which is also `grpadmin`. Note that passwords are *not* echoed on the console.

```
Login: grpadmin
Password:
```

```
        Welcome to Group Manager
        Copyright 2001 - 2007 EqualLogic, Inc.
```

```
It appears that the storage array has not been configured.
Would you like to configure the array now? (y/n) [n]
```

4. When prompted, enter `y` to invoke the `setup` utility. You can also run `setup` later at the console prompt (`>`). At each `setup` prompt, enter a question mark (`?`) for help. Press the `[Enter]` key to accept a default (shown in brackets).
5. Specify the member configuration and group information at the prompts. This is the information described in Table 2-2. In addition, you will be prompted for the name, IP address and membership password for the group you want to join. There may be a short delay as `setup` locates the group on the network.

Once the array is successfully added to the new group, use the Group Manager GUI or CLI to set the RAID policy for the new member.

The storage in the new member will be available once you set the RAID policy. The disks are automatically configured according to the designated RAID level, with the appropriate number of spare disks. Until the RAID configuration completes, performance will not be optimal, but the group is fully operational.

See *Using the Group Manager Graphical User Interface* on page 2-7 and *Using the Group Manager Command Line Interface* on page 2-19 for information on the starting the Group Manager user interfaces. See *Introduction to Member RAID Policies* on page 9-1 for information on setting the RAID policy.

Using the Group Manager Graphical User Interface

The Group Manager graphical user interface (GUI) enables you to perform group administration tasks using windows and dialog boxes.

There are two methods you can use to run the Group Manager GUI:

- Use a Web browser through a standard Web connection using HTTP (port 80).
- Install the GUI on a local system and run it as a standalone application, as described in *Running the GUI on a Local System* on page 2-17.

Notes: See the *Release Notes* for information about browser requirements for running the GUI. If the window size is not sufficient, a pop-up message will appear.

By default, Web access to the group is enabled. For information about enabling and disabling access, see *Controlling GUI and CLI Management Access* on page 3-7.

To run the Group Manager GUI on a Web browser, specify the group IP address in a Web browser window, as follows:

```
http://group_ip_address
```


When you connect to the group, the login dialog box appears (Figure 2-1). Enter a group administration account name. You can use the default `grpadmin` account and password you set when creating the group if no other accounts have been set up. For information about setting up accounts, see *Managing Group Administration Accounts* on page 3-9.

Figure 2-1: Group Manager Login



After you log in to the group, the Group Summary window appears and displays an overview of the group status, as shown in Figure 2-2. If you have not chosen the RAID policy for the member, this condition will be indicated by a message in the Storage Pools and Group Members panel.

Figure 2-2: Group Summary



The GUI enables you to manage various group components, which you access by clicking and expanding objects in the tree structure in the far left panel:

- **Group Configuration** – Modify the group configuration and set up accounts, event notification, network services, authentication, and SNMP. See Chapter 3, *Group Customization*.
- **Group Monitoring** – Monitor iSCSI connections to the group, snapshot and replication schedules, volume replication configurations and activity, administrative sessions and login history, and in-progress member and volume move operations. See *Monitoring the Group* on page 10-7.
- **Events** – Display events in the group. See *Monitoring Events* on page 10-1.
- **Storage Pools** – Create and manage pools in the group. See Chapter 4, *Storage Pool Management*.
- **Members** – Monitor and manage group members, including configuring network interfaces. See Chapter 9, *Member Management*.
- **Volumes** – Monitor and manage volumes, snapshots, replicas, and schedules. See Chapter 5, *Volume Management*.
- **Volume Collections** – Create and manage collections of volumes. Organizing multiple, related volumes into a collection enables you to create snapshots or replicas of the volumes in a single operation or schedule. See *Managing Volume Collections* on page 5-15.
- **Replication Partners** – Monitor and manage replication partners. See Chapter 7, *Basic Replication*.

As you expand and select objects in the tree, progressively more detailed information appears in the far right panel. Most main windows include an Activities panel that displays context-specific tasks for the selected object.

Icons at the top of each window enable you to do the actions described in Table 2-4.

Table 2-4: GUI Icons

	Apply changes. If you enter changes in a GUI window, click this icon to apply the changes. If you do not apply the changes, you will be prompted to do so when you leave the window.
	Discard changes. If you make a mistake and want to discard changes you made in a window, click this icon.
	Refresh data in a window. Click this to refresh the GUI data. Do not use the browser refresh button to refresh the data in the GUI.
	Navigate the GUI. Click one of these icons to move backward or forward through the GUI windows, according to the window history.

The Outstanding Alarms panel at the bottom of the GUI window displays a list of alarms in the group, so you can quickly discover when problems occur, in addition to a list of user tasks that are in progress:

- Click the arrow in the Outstanding Alarms title bar or click the title bar to display and hide the panel. Each alarm contains a link to the affected object (member or volume). Click the link for additional information.
- Click the double checkmark to acknowledge all alarms and turn off the flashing yellow information icon.

Click the Tools title bar or the arrow at the far left bottom of the GUI window to display or hide the Tools panel, which contains the following links:

- `User preferences` – Control aspects of GUI behavior. See *Setting GUI User Preferences* on page 2-11.
- `Online help` – Display topic-oriented help on every aspect of using the GUI and the CLI to manage a group. See *Obtaining Online Help* on page 2-16.
- `Customer support` – Launch the Technical Support website. See *Technical Support and Customer Service* on page xiii.
- `Performance monitor` – Display various performance statistics for the group. See *Monitoring Performance Statistics* on page 10-15.
- `Manual transfer utility` – Launch the Manual Transfer Utility. See *Displaying Replication Status* on page 7-26.

- **Run as application** – Install the Group Manager GUI on the system that is running the GUI and run Group Manager as a standalone application. If the GUI is already installed locally, clicking the link will launch the application. See *Running the GUI on a Local System* on page 2-17.

Controlling Group Manager GUI Operation

You can control how the GUI operates and displays information as follows:

- *Securing GUI Access* on page 2-11
- *Setting GUI User Preferences* on page 2-11
- *Obtaining Online Help* on page 2-16
- *Running the GUI on a Local System* on page 2-17

Securing GUI Access

To secure communication between your Web browser and the group, select the checkbox next to `Encrypt communication` in the Group Manager Login dialog box (Figure 2-1).

If the checkbox is grayed out, you must download and install the latest version of the Java™ Runtime Environment (JRE) to take advantage of Java security features. To download and install the Java plug-in, go to the `www.java.com` website. Encrypted communications are supported on Java Runtime Environment 1.4 and higher versions.

To *require* encryption between the Web browser and the group, you must also click:

`Group Configuration > Administration tab`

In the Group Administration window that appears (Figure 3-4), select `Allow only secure connections`.

Setting GUI User Preferences

To control the GUI display and behavior, you can set the following:

- *General Policies* on page 2-12
- *Communication Policies* on page 2-13
- *Alarm Policies* on page 2-14
- *Data Validation and Debugging Policies* on page 2-15


General Policies

To set general GUI policies, click:

Tools > User Preferences > General tab

The User Preferences – General dialog box appears (Figure 2-3).

Figure 2-3: User Preferences – General



You can set the following general GUI policies:

- GUI font size.
- Location of the Activities panel, in relation to the main GUI panel.
- Location of the Group Manager online help. For more information, see *Obtaining Online Help* on page 2-16.
- Whether the GUI displays tooltips and popup descriptions for list items.

Select the general settings and click OK.


Communication Policies

To set GUI communication policies, click:

Tools > User Preferences > Communication tab

The User Preferences – Communication dialog box appears (Figure 2-4).

Figure 2-4: User Preferences – Communication



The User Preferences – Communication dialog box allows you to set:

- **Connection policies.** Control whether the GUI will attempt to reconnect with the group when the TCP/IP connection between the browser and the group is lost.

If you select `Automatically reconnect if disconnected`, the GUI will automatically attempt to reconnect for 100 seconds. If the browser fails to reconnect, it will display a dialog box informing you that you have lost your connection and asking if you still want to reconnect. If you answer `yes`, the GUI will make another attempt. If you answer `no`, a message appears in red in the upper lefthand corner stating that the connection has been lost. If you de-select the policy, the GUI makes no automatic attempt at reconnection and the dialog box appears immediately.

If you de-select **Keep connection alive** when temporarily leaving GUI page, you will lose your connection if you click the **Back** (or **Forward**) button in your browser. If you select the policy, the connection is kept alive for 10 minutes.

- **Maximum data refresh interval.** Specifies the maximum time interval between GUI data updates. The default setting is 60 seconds. For example, using the default value, GUI data will be updated every 60 seconds or less, depending on the group workload. The minimum interval is 5 seconds.

To manually update data displayed in the GUI, click the **Refresh data** icon in the upper righthand corner of the GUI (under the **Logout** button). Do not use the browser's refresh button.

Select the desired policies and click **OK**.


Alarm Policies

To control the behavior of alarm notification and the event log, click:

Tools > User Preferences > Alarms tab

The User Preferences – Alarms dialog box appears (Figure 2-5).

Figure 2-5: User Preferences – Alarms



The dialog box enables you to customize:

- **Failure notification.** You can control how the GUI indicates that a failure has occurred. You can choose to have the GUI flash the status icons of the failed objects or highlight failed objects using the color red.
- **New alarm notification.** When new alarms occur, you can choose to have the GUI open pop-up windows for them, or emit an audible sound.
- **Event logging.** You can choose how events are displayed when you click **Events** in the far left panel. You can choose to display the events in a separate window or choose to automatically retrieve the 100 most recent events when you start the GUI.

Select the desired policies and click **OK**.


Data Validation and Debugging Policies

To control the level of validation of user input and whether debugging information is logged to a Java console, click:

Tools > User Preferences > Advanced tab

The Modify User Preferences Advanced dialog box appears (Figure 2-6).

Figure 2-6: Modify User Preferences Advanced



You can customize advanced user preferences as follows:

- **Data validation behavior.** You can choose whether text is checked for errors as you type it or when the information is submitted.
- **Debug logging.** You can choose to log data poller activities and server notifications to a Java console. If enabled, these options can slow performance, so keep them disabled for better performance.

Obtaining Online Help

When you expand `Tools` and click `Online Help` in the far left panel in the GUI, the Group Manager online help appears in a separate window, providing you with topic-oriented help on every aspect of managing a group using the GUI and CLI.

By default, the online help is launched from a directory that reflects the PS Series firmware version, such as the following:

```
www.psonlinehelp.equallogic.com/V3.3
```

The Group Manager online help files are also located on the CD-ROM shipped with each PS Series array.

If desired, you can change the default online help location to a local network. Follow these steps:

1. Add the group IP address to the list of trusted sites, if necessary. Depending on the version of Java running on the local system (for example, Java 1.4.2), to access the online help from a local directory, you may need to add the group IP address to the system's list of trusted sites in Internet Explorer:
 - a. Launch Internet Explorer and under `Tools`, select `Internet Options`.
 - b. Select the `Security` tab and click `Trusted Sites` and then the `Sites` button.
 - c. In the `Trusted Sites` window, enter the group IP address in the `Add this Web site to the zone` field using this format:

```
http://group_ip_address
```

Do not select `Require server verification (https:)` for all sites in this zone, which appears at the bottom of the `Trusted Sites` window.

- d. Click `Add`.

2. Copy the `groupmanagerhelp` directory from the product CD-ROM to a network drive accessible to all systems that will be using the GUI online help.
3. Open the GUI (enter the group IP address in a browser) and click:

Tools > User Preferences > General tab

The User Preferences General dialog box appears (Figure 2-3). In the `Location of online help files` field, enter the help directory location and click **OK**. For example, to enter a network location:

```
file:///system.directory.company.com/share/groupmanagerhelp
```

You can also put the online help files on a computer connected to your network, navigate to that directory using a browser, and enter `groupmanager.htm` in the browser's address field.


Running the GUI on a Local System

You can install the Group Manager GUI on a local system and run the GUI as a standalone application. This provides a fast, convenient way to start the GUI, uses less system memory, provides more usable window space, and eliminates browser security issues. If desired, you can have GUIs for more than one group installed on a system.

Note: A system must be running Sun Java version 1.4 or a later version to run the GUI as a standalone application.

Follow these steps to install the GUI on a local system:

1. On the system on which you want to install the GUI, specify the group IP address in a Web browser window and log in to a group administration account (for example, the `grpadmin` account).
2. In the Group Summary window (Figure 2-2), expand **Tools** in the GUI and click **Run as application**.
3. When prompted to install the GUI on the local system to manage the current group, click **Yes**.
4. In the Create Shortcut dialog box (Figure 2-7), choose whether to create a shortcut on the desktop.

Figure 2-7: Create Shortcut

If you click **Yes**, the GUI will be launched locally. In addition, a shortcut will be installed on your desktop for managing the current group, and the GUI will appear in the Startup menu of Microsoft Windows systems.

If you click **No**, the GUI will be launched locally, but no shortcut will be created. You can launch the GUI by using the Java Control Panel or Java Web Start, depending on your version of Java. For example, using the Java Control Panel (available in the Windows Control Panel), click the **General** tab, **Settings**, and **View Applications** and then double-click the applet.

If you click **Configure**, the Java Control Menu will appear, enabling you to modify the Java settings.

Once the GUI is installed locally, it will be automatically updated when you update the PS Series firmware.

Note: If you change the IP address of a group for which you are running the GUI locally, you must uninstall the GUI from the system and then reinstall it.

To uninstall a standalone GUI application, you can use the Java Control Panel or Java Web Start. For example, in the Java Control Panel, click the **General** tab, **Settings**, and **View Applications**. The GUI application for a group has two components, the applet and the library; both are prefaced with the group name. Be sure to delete both components.

Using the Group Manager Command Line Interface

The Group Manager CLI enables you to enter commands on a command line (or in a host-based script) to manage a group or a specific array.

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

- **Use a network connection.** From a host, use telnet or SSH to connect to the group 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 serial port 0 on the active control module (the ACT LED will be green). Use the serial cable that shipped with the array. The serial connection must have the following characteristics:
 - 9600 baud
 - One STOP bit
 - No parity
 - 8 data bits
 - No hardware flow control
- Once connected to the group, log in to a group administration account, such as `grpadmin`. You can then enter commands at the CLI prompt.

There are three types of CLI commands:

- **Group management commands.** These commands manage the group, members, and volumes. They use a tree structure and generally require subcommands, parameters, or variables.
- **Array management commands.** These commands perform tasks *only* on a specific array. To run, you must use a telnet or ssh connection to an IP address assigned to a network interface on the array, *not* to the group IP address. Alternately, you can use a serial connection to the array.
- **Global commands.** These commands provide access to common command options and can be executed anywhere in the CLI tree structure.

By default, telnet and SSH access to the group is enabled. For information about enabling and disabling access, see *Controlling GUI and CLI Management Access* on page 3-7.

For detailed information about using the CLI to manage a group and individual arrays, see the PS Series *CLI Reference* manual.

Where to Go Next

After setting up a group, you may want to customize the group or implement advanced group features. With the exception of installing disks, use the Group Manager graphical user interface (GUI) or command line interface (CLI) to perform the operations described in the following tables.

Table 2-5: Group Administration and Event Notification

To:	Do this:	Reference:
Modify the group date, time, and time zone	By default, the time zone is EST (eastern standard time). Daylight saving time is applied if it is used in the time zone.	<i>Modifying the Date, Time, and Time Zone</i> on page 3-6
Control who can perform group administration tasks	Create accounts to restrict the level of access for administrators and to monitor their activity.	<i>Managing Group Administration Accounts</i> on page 3-9
Be notified when events occur in the group	Set up e-mail or syslog event notification.	<i>Setting Up Event Notification</i> on page 3-20
Monitor traps from the group	Configure SNMP.	<i>Using SNMP to Monitor a Group</i> on page 3-31

Table 2-6: Network Services

To:	Do this:	Reference:
Automate iSCSI target discovery	Specify an iSNS server for the group to use.	<i>Using iSNS for iSCSI Device Management</i> on page 3-30
Synchronize group time	Specify an NTP server for the group to use.	<i>Modifying the Date, Time, and Time Zone</i> on page 3-6

Table 2-7: Volumes

To:	Do this:	Reference:
Provide disk space to users or applications	Create one or more volumes.	<i>Creating Volumes</i> on page 5-3
Restrict host access to volumes	Set up one or more access control records for a volume.	<i>Managing Access Controls for Volumes and Snapshots</i> on page 5-23

Table 2-7: Volumes (Continued)

To:	Do this:	Reference:
Use CHAP for initiator authentication	Set up local CHAP accounts in the group or specify an external RADIUS server.	<i>Using CHAP to Control iSCSI Access to Volumes</i> on page 3-25
Mutually authenticate iSCSI targets and initiators	Set up target authentication, in addition to initiator authentication.	<i>Configuring Target Authentication</i> on page 3-29
Access a volume (or snapshot) from a host	Install an iSCSI initiator of your choice on a host and connect to a volume (or snapshot).	<i>Connecting Hosts to Volumes and Snapshots</i> on page 5-27
Capture the current contents of a volume	Create a snapshot. You can clone a snapshot and restore a volume from a snapshot.	Chapter 6, <i>Snapshot Management</i>
Create a new volume that is a copy of the original	Clone a volume.	<i>Cloning Volumes</i> on page 5-38
Protect volume data from catastrophic failure or the destruction of a volume	Set up volume replication across two groups.	Chapter 7, <i>Basic Replication</i>
Group volumes for the purpose of creating snapshots or replicas of the volumes in a single operation or schedule	Create a volume collection.	<i>Managing Volume Collections</i> on page 5-15
Restrict volume data to members with a specific RAID level	Set the preferred RAID level for a volume.	<i>Setting a Volume RAID Preference</i> on page 5-35

Table 2-8: Capacity and Performance

To:	Do this:	Reference:
Increase network bandwidth and improve availability and performance	Configure additional network interfaces (up to three) on a group member.	<i>Managing the Member Network Configuration</i> on page 9-18
	Add members to the group.	<i>Adding Members to an Existing Group</i> on page 2-5

Table 2-8: Capacity and Performance (Continued)

To:	Do this:	Reference:
Increase group storage capacity	Install disks in a group member.	<i>Installing and Removing Disks</i> on page 9-16
	Add members to the group.	<i>Adding Members to an Existing Group</i> on page 2-5
Divide group space for the purpose of dedicating storage to different workloads	Create multiple storage pools and assign members and volumes to the pools.	Chapter 4, <i>Storage Pool Management</i>

3 Group Customization

Once you create a PS Series group, you can customize the group.

Only group administrators can customize a group. For more information, see *Managing Group Administration Accounts* on page 3-9.

Group customization topics include:

- *Displaying the Group Configuration and Status* on page 3-1
- *Modifying the Group Network Configuration* on page 3-5
- *Modifying the Date, Time, and Time Zone* on page 3-6
- *Controlling GUI and CLI Management Access* on page 3-7
- *Managing Group Administration Accounts* on page 3-9
- *Setting Up Event Notification* on page 3-20
- *Using CHAP to Control iSCSI Access to Volumes* on page 3-25
- *Using iSNS for iSCSI Device Management* on page 3-30
- *Using SNMP to Monitor a Group* on page 3-31
- *Modifying Access Controls for Windows Services* on page 3-35
- *Modifying the Group-Wide Volume Settings* on page 3-36
- *Controlling Performance Load Balancing* on page 3-38
- *Preventing Member Firmware Downgrades* on page 3-40

Displaying the Group Configuration and Status

When you first log in to the GUI, the Group Summary window (Figure 3-1) appears. This window displays an overview of the group configuration and status. Any alarms are shown in the panel at the bottom of the window.

Figure 3-1: Group Summary



The Group Information panel shows the number of volumes, snapshots, and collections in the group, in addition to the number of volumes and snapshots that are online and in use (that is, have a current iSCSI connection).

Note: A volume can be offline due to a failure or administrator action. In addition, a thin provisioned volume can go offline if its volume reserve reaches the maximum in-use limit. If volumes are offline, you should investigate whether a failure has occurred. See *Displaying Volume Status and Configuration* on page 5-11 for information about monitoring volumes. By default, snapshots are offline, although you can manually set a snapshot online.

In the Group Disk Space panel, select the type of information you want to display:

- Group space utilization – Group space allocated to volumes, snapshots, local replication reserve, and delegated space and the percentage of free group space.

If free group space is low, you can add members to the group, as described in *Adding Members to an Existing Group* on page 2-5.

- RAID level space distribution – Group space configured with each RAID level (RAID 10, RAID 5, and RAID 50).
- Storage pool capacity – Percentage of group space in each storage pool and the capacity of each pool.
- Delegated space utilization – Used and free group space delegated to all replication partners.

If free delegated space is low, monitor the replication partners, as described in *Displaying Replication Activity Between Partners* on page 7-32, to determine which partner needs more delegated space. See *Modifying a Replication Partner* on page 7-17 for information on increasing delegated space.

Note: The group free space shown in the GUI may not reflect the actual free space because the GUI uses a rounding algorithm to calculate free space. To display the actual free space in megabytes, use the CLI command `cli-settings displayinMB on`. Then use the `show` command to display the free group space.

The Storage Pools and Group Members panel shows the members, capacity, and the percentage of free space in each storage pool.

A member that has a warning health condition will appear in orange with the warning symbol. A member with a critical health condition will appear in red with an error symbol. An offline member appears in gray. See *Displaying Member Configuration and Status* on page 9-6 for information about monitoring members and diagnosing warning and critical health conditions.


Note: If a pool's free space is low, consider moving volumes from the low-space pool to a different pool, as described in *Moving a Volume Across Pools* on page 5-29. You can also move a member to the low-space pool, as described in *Moving a Member Across Pools* on page 9-27.

To customize the group, click:

Group Configuration

The Group General window appears (Figure 3-2).

Figure 3-2: Group General



Click the window tabs to display group information and customize the group:

- **General** – Modify the group name, IP address, location, and description, in addition to the date, time, time zone, and Network Time Protocol (NTP) server configuration for the group. See *Modifying the Group Network Configuration* on page 3-5 and *Modifying the Date, Time, and Time Zone* on page 3-6 for more information.
- **Administration** – Manage group administration accounts, both locally and using RADIUS authentication and accounting, and control how users can access a group. See *Controlling GUI and CLI Management Access* on page 3-7 and *Managing Group Administration Accounts* on page 3-9.
- **Notifications** – Manage group event logging and event notification methods. See *Setting Up Event Notification* on page 3-20.
- **iSCSI** – Configure initiator and target authentication for iSCSI access to volumes and snapshots and configure the group to use iSNS servers. See *Using CHAP to Control iSCSI Access to Volumes* on page 3-25 and *Using iSNS for iSCSI Device Management* on page 3-30.
- **SNMP** – Manage SNMP access to the group for monitoring purposes. See *Using SNMP to Monitor a Group* on page 3-31.

- **VDS/VSS** – Modify the access control records that allow Microsoft® Windows® services to access to the group for management operations. See *Modifying Access Controls for Windows Services* on page 3-35.
- **Defaults** – Modify the default settings for all volumes in the group (snapshot space, thin provisioning, and iSCSI alias settings). See *Modifying the Group-Wide Volume Settings* on page 3-36.
- **Advanced** – Specify the performance load balancing setting and disallow firmware downgrades. See *Controlling Performance Load Balancing* on page 3-38 and *Preventing Member Firmware Downgrades* on page 3-40.

For details about interpreting group status, see *Displaying the Group Configuration and Status* on page 3-1.

Modifying the Group Network Configuration

To modify the group network configuration, including the group name, IP address, location, and description, click:

Group Configuration > General tab

The Group General window appears (Figure 3-2). In the General Settings panel, you can modify the following fields:

- **Group Name** – Name used to identify a group for administrative purposes and for volume replication. See Table 2-2 on page 2-3 for group name requirements.
- **Group IP address** – Network address used for group administration (GUI and CLI) and for host access to volumes. Be sure to follow the network requirements described in Table 2-1 on page 2-2.

Note: Modifying the group IP address will disconnect current group administration and iSCSI connections to the group. You must inform administrators and hosts of the new group IP address.

- **Location** – Location (physical) of the group. A location is optional.
- **Description** – Description of the group. A description is optional.

Caution: If you modify the group name or IP address and the group is configured as a replication partner on another group, be sure to log in to the other group and make the reciprocal name or IP address change. Replication will fail if the partner information is incorrect. See *Configuring Replication Partners* on page 7-11 for information.

Modifying the Date, Time, and Time Zone

You can modify the group date, time, and time zone, as needed. You can also specify an external Network Time Protocol (NTP) server that the group will use to automatically synchronize time on all the members.

A group ensures that all members use a consistent time. The clock in each array is set at the factory and is based on GMT. The default time zone is America/New York. The group time is based on the first member's clock. When an array joins a group, its clock is updated to the group date and time.

If the group time changes (either because you manually modified it or because the NTP server time changed), the time is first set on one group member. Then, the other group members synchronize their times with this member. Synchronizing time on a member can take as long as 24 hours if the time changes by an hour and could take longer if the time changes by more than an hour.


To modify the group date, time, or time zone or to specify an NTP server, click:

Group Configuration > General tab

The Group General window appears (Figure 3-2).

In the Date and Time panel, you can modify:

- **Time zone.** Select a time zone from the `Time zone` drop-down menu. Whether daylight saving time is applied depends on the time zone.
- **Date or time.** Click `Modify group time`. The `Modify Group Time` dialog box appears (Figure 3-3). Modify the date and time, as needed, and click `OK`. You can also click `Set to time of system running GUI` to set the time to match the time on the system running the GUI. Click `Reset to current group time` to cancel the changes.

Figure 3-3: Modify Group Time

Alternatively, you can obtain the group time from an NTP server.

In the NTP servers panel, click **Add**. In the Add list item dialog box, specify the IP address for an NTP server and click **OK**. You can specify up to three IP addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows to change the address order.

Use the *ip_address:port* format to specify a port number if you are using a port other than the default, 123.

As needed, you can modify or delete an NTP server by selecting the server IP address and clicking **Modify** or **Delete**.

Note: An NTP server must be set up and available on the network in order for the group to access it.

Controlling GUI and CLI Management Access

You can control whether administrators can use the Group Manager GUI or CLI to manage a group. You can also specify restrictions on those access methods.

Note: Read-only access to the group is available through SNMP if you set the community string. See *Using SNMP to Monitor a Group* on page 3-31 for more information. If you do not have network access to the group, you can use a serial connection to access the group and run CLI commands. See *Using the Group Manager Command Line Interface* on page 2-19 for more information.

To control GUI and CLI access to the group, click:

Group Configuration > Administration tab

The Group Administration window appears (Figure 3-4).

Figure 3-4: Group Administration



You can control the following access methods in the Administration Access panel:

- **GUI access.** To enable use of the GUI to manage the group, select the checkbox next to Enable Web access. De-select the checkbox to disallow all use of the GUI.

To require encryption between a Web browser and the group, also select the checkbox next to Allow only secure SSL connections. In addition, you must download and install Java Runtime Environment (JRE), which provides SSL, and select the checkbox next to Encrypt communication in the Group Manager login dialog box (Figure 2-1).

- **CLI access.** To use the CLI to manage the group, select Enable telnet access or Enable SSH (secure shell) access, as desired. De-select both checkboxes to disallow use of the CLI.

Managing Group Administration Accounts

To use the Group Manager GUI or CLI, you need to log into a group administration account. Administration accounts prevent unauthorized individuals from accessing a group.

The default administration account, `grpadmin`, allows you to perform all group operations. It is recommended that you set up a separate account for each administrator to track their activities.

There are three account types, with different permissions:

- **Group administrator.** This account type has the same permissions as the default `grpadmin` account to modify any and all aspects of the group and all pools. Only group administrators can create or modify accounts.
- **Pool administrator.** This account type can perform operations only on the objects (volumes, members, snapshots, and so on) in the pool or pools to which the account has permission. However, pool administrators cannot move objects between pools they manage. Pool administrators can optionally monitor the entire group.
- **Read-only account.** This type of account cannot modify any objects in the group, but can view group and pool information. Use this account type for administrators who will only monitor group activity.

You can set up administration accounts in two ways:

- **Local accounts** stored and managed in the group. See *Using Local Administration Accounts* on page 3-9 for more information.
- **External accounts** stored and managed on a RADIUS server. You can also use accounting to track the activity of these accounts. For more information, see *Using External RADIUS Accounts* on page 3-13.

Using Local Administration Accounts


To manage local group administration accounts, click:

Group Configuration > Administration tab

The Group Administration window appears (Figure 3-4). The Administration Accounts panel displays the default `grpadmin` account and any local accounts that you have created.

To create a local account, click **Add**. The **Create Account – General Settings** dialog box appears (Figure 3-5).


Figure 3-5: Create Account - General Settings



Specify information in the following fields:

- **Account 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. Only the first 8 characters are validated.
- **Description** — Description of the account (optional).

After entering the general settings, click **Next**. The **Create Account – Account Permissions** dialog box appears (Figure 3-6).


Figure 3-6: Create Account - Account Permissions

Select the type of account.

For pool administrator accounts, also select the pool or pools that the account can manage. Optionally, you can allow a pool administrator to monitor the entire group by clicking `Read access to entire group`.

By default, the checkbox for `Enable administration account` is selected. De-select the checkbox to disable the account and prevent someone from using it to log in to the group.

After specifying the account type, click `Next`. The `Create Account – Contact Information` dialog box appears (Figure 3-7).

Figure 3-7: Create Account - Contact Information

Create group administration account

Step 3 - Contact Information

Contact information:

Contact name: Kristen

E-mail address: kristen@company.com

Phone:


Mobile:

[< Back](#) [Next >](#) [✓ Finish](#) [✗ Cancel](#)

Java Application Window

After entering the optional contact information, click **Next**. The Create Account – Summary dialog box (Figure 3-8) appears, displaying a summary of the account information.

If the information is satisfactory, click **Finish** to create the account.

Figure 3-8: Create Account - Summary

Create group administration account

Step 4 - Summary

Summary of account settings:

General Settings

Account name: kristen

Account type: pool administrator (group-read-access)

Managed pools: pool2

Account status:  enabled

Description:

Contact Information

Name: Kristen

E-mail: kristen@company.com

Phone:

Mobile:

[< Back](#) [Next >](#) [✓ Finish](#) [✗ Cancel](#)

Java Application Window

Group administrators can modify accounts in the following ways:

- Change any account type to any other account type. (However, you cannot change the account type for the default group administrator account, `grpadmin`).
- Add or delete pools from a pool administrator account.

Notes: You cannot modify the name of an administration account. Instead, you must delete the account and then add it again with the new name.

You cannot disable or delete the `grpadmin` account.

If you change an account type and are using an external RADIUS authentication server, make sure to change the corresponding vendor-specific attribute for the account type. See Table 3-1 for more information.

To modify or delete an account, in the Group Administration window (Figure 3-4), select the account and click **Modify** or **Delete**. Change the account information or confirm that you want to delete the account, as desired.

Using External RADIUS Accounts

Using an external RADIUS authentication server to manage group administration accounts can be useful when you have many accounts (more than 100). You can also use RADIUS accounting servers to track the activity of these accounts.

RADIUS servers must be available on the network and accessible to the group for the servers to authenticate the account login or perform accounting.

There are various implementations of RADIUS. For information about setting up RADIUS servers themselves and specifying vendor-specific attributes (VSAs), see the documentation for your product.

Note: RADIUS authentication servers can also be used to authenticate host iSCSI access to the group. See *Using External RADIUS Servers for Initiator Authentication* on page 3-28.

Setting Up External RADIUS Authentication Servers

Setting up an external RADIUS authentication server and configuring a group to use it involves a series of steps that alternate between those that you perform on the group, and those that you perform on the RADIUS server.

Note: Pool administrator accounts defined on a RADIUS server will be allowed full administrator access to PS Series groups running firmware versions 3.0.x and 3.1.x unless you configure the RADIUS server to restrict access, based on the IP address, to only groups running 3.2 or greater. It is recommended that you restrict access based on the client source IP address.

To configure a group to use an external RADIUS authentication server:

1. On the group, click:

Group Configuration > Administration tab

The Group Administration window appears (Figure 3-4).

2. Select the checkbox next to `Enable RADIUS Authentication for login`.

By default, the `Require vendor-specific RADIUS attribute` option is selected. This means that the `EQL-Admin` attribute must be configured as part of an account on the RADIUS server, and the RADIUS server must return it to the group in the Access-Accept message. For security reasons, it is strongly recommended that you require the `EQL-Admin` attribute.

Optionally, select `Enable RADIUS accounting for authenticated users` if you want to track account activity. You will set up the RADIUS accounting server in a later step.

3. On the RADIUS server, you must specify one of the following RADIUS attributes:

- `Service-Type = Administrative` (for read-write access)
- `Service-Type = NAS-Prompt` (for read-only access)

If the `Service-Type` is not set to either value, the administrator login will fail. See your RADIUS documentation for information on how to set attributes.

4. If you selected the `Require vendor-specific RADIUS attribute` option in Step 2, you must configure the vendor-specific attributes (VSAs) for each account, using the information in Table 3-1. If present on the RADIUS server, these attributes will be

displayed (except for EQL-Admin-Poll-Interval) in the Contact Information fields in the Group Manager GUI.

Table 3-1: Vendor-Specific Attributes

Attribute	Field	Required Value
EQL-Admin Defines the type of login account. Required if the Require vendor-specific RADIUS attribute option is selected in the GUI.	VSA Vendor ID	12740
	VSA Type	6
	VSA Length	6 (1 byte type + 1 byte length + 4 bytes value)
	VSA Syntax	integer (4 bytes)
	VSA Value	1 (group administrator) 2 (pool administrator)
EQL-Pool-Access (Required only if the value of the EQL-Admin attribute is 2, for pool administrator accounts. If so, use the VSA Syntax field to list the pools to which the pool administrator has access.	VSA Vendor ID	12740
	VSA Type	7
	VSA Length	3 - 255
	VSA Syntax	string (comma-separated list of pools)
EQL-Admin-Full-Name (Optional). The name of the administrator who will use the account.	VSA Vendor ID	12740
	VSA Type	1
	VSA Length	3-253
	VSA Syntax	string
EQL-Admin-Email (Optional). The email address of the administrator.	VSA Vendor ID	12740
	VSA Type	2
	VSA Length	3-253
	VSA Syntax	string
EQL-Admin-Phone (Optional). The phone number for the administrator.	VSA Vendor ID	12740
	VSA Type	3
	VSA Length	3 - 253
	VSA Syntax	string

Table 3-1: Vendor-Specific Attributes (Continued)

Attribute	Field	Required Value
EQL-Admin-Mobile (Optional). The mobile phone number for the administrator.	VSA Vendor ID	12740
	VSA Type	4
	VSA Length	3 - 253
	VSA Syntax	string
EQL-Admin-Poll-Interval (Optional). Specifies how often, in seconds, the group configuration data must be re-pollled by the GUI. Default is 30.	VSA Vendor ID	12740
	VSA Type	5
	VSA Length	6
	VSA Syntax	integer


5. On the group, specify the IP address for the RADIUS authentication server, as described in *Configuring RADIUS Authentication Servers* on page 3-16.
6. Optionally, specify the IP address for a RADIUS accounting server if you want to use accounting on the accounts, as described in *Configuring RADIUS Accounting Servers* on page 3-18.

Configuring RADIUS Authentication Servers

Before configuring RADIUS authentication servers in the Group Manager GUI, you must set up the RADIUS servers as described in *Setting Up External RADIUS Authentication Servers* on page 3-13.


To configure a RADIUS authentication server in the group, click **RADIUS settings**. The RADIUS Settings dialog box appears (Figure 3-9).

Figure 3-9: RADIUS Settings



Under RADIUS authentication servers, click **Add**. The Add RADIUS Authentication Server dialog box appears (Figure 3-10).

Figure 3-10: Add RADIUS Authentication Server



Specify the IP address for the server. Use the format *ip_address:port* if the port is different from 1812.

Optionally, enter the secret (password) for the RAIDUS authentication server. The secret can contain up to 63 characters. It is recommended that you use secrets for security reasons.

Click **OK** when you are finished.

You can specify up to three IP addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows in the RADIUS Settings dialog box to rearrange the addresses.

To modify or delete an IP address, in the RADIUS Settings dialog box, select the IP address and click **Modify** or **Delete**.

In addition, for all the RADIUS authentication servers specified in the RADIUS Settings dialog box (Figure 3-9), specify information in the following fields:

- **Request timeout** – Number of seconds the group will wait for an authentication server to respond before timing out. The default is 2 seconds.

For example, when a user attempts to log into an administration account, the group will try to contact an authentication server. If the server does not respond within the request timeout period, the attempt at contact will fail. Depending on the retry setting, the group may again try to contact the server.


- **Number of retries** – Number of times the group will try to contact an authentication server after the first attempt fails. The default is 1 retry.

Once the retry limit has been reached, the group will consider the server unavailable. If there is another authentication server configured, the group will then attempt to contact that server.

For example, when a user attempts to log into an administration account, the group will try to contact an authentication server. If there is no response within the request timeout period, the group will retry contacting the server the number of times specified by the retry setting.

Configuring RADIUS Accounting Servers

To configure a RADIUS accounting server, in the RADIUS Settings dialog box (Figure 3-9), under RADIUS accounting servers, click **Add**. The Add RADIUS Accounting Servers dialog box (Figure 3-11) appears.

Figure 3-11: Add RADIUS Accounting Server

Add RADIUS accounting server

Enter IP address of RADIUS accounting server.
Specify the port number if it is different than 1813
(use format <ip_address>:<port_number>):

IP address: 192.12.44.120

RADIUS secret: *****

Confirm secret: *****

 OK  Cancel

Java Application Window

Specify the IP address for the server. Use the format *ip_address:port* if the port is different from 1813.

Optionally, enter the secret (password) for the RAIDUS authentication server. The secret can contain up to 63 characters. It is recommended that you use secrets for security reasons. Click **OK** when you are finished.

You can specify up to three IP addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows in the RADIUS Settings dialog box to rearrange the addresses.

To modify or delete an IP address, in the RADIUS Settings dialog box, select the IP address and click **Modify** or **Delete**.

In addition, for all the RADIUS accounting servers specified in the RADIUS Settings dialog box (Figure 3-9), specify information in the following fields:

- **Request timeout** — Number of seconds the group will wait for an accounting server to respond before timing out. The default is 2 seconds.

For example, when a user attempts to log into an administration account, the group will try to contact a RADIUS authentication server. Then—if the account is authenticated—the group will try to contact a RADIUS accounting server. If the accounting server does not respond within the request timeout period, the attempt at contact will fail. Depending on the retry setting, the group may again try to contact the accounting server.

- **Number of retries** – Number of times the group will try to contact an accounting server after the first attempt fails. The default is 1 retry.

For example, when a user attempts to log into a group administration account, the group will try to contact a RADIUS authentication server. Then—if the account is authenticated—the group will try to contact an accounting server. If there is no response within the request timeout period, the group will retry contacting the server the number of times specified by the retry setting.

Setting Up Event Notification

It is recommended that you set up event notification to be informed automatically when events occur in the group.

If an event occurs in the group (such as the creation of a volume or the removal of a power supply), a message is generated. You can use these messages to track normal operations and also to detect and diagnose problems before they disrupt operations. View event messages by clicking **Events** in the far left GUI panel.

Each event is assigned a priority. Table 3-2 lists event priorities in order from highest (most severe) to lowest (least severe).

Table 3-2: Event Priorities

Priority	Description
FATAL	Catastrophic failure that requires immediate administrator intervention.
ERROR	Serious failure that should be investigated and may require administrator intervention.
WARNING	Potential problem in the group that may become an event with ERROR priority if administrator intervention does not occur.
INFO	Normal operational and transitional event that is of interest but does not require intervention.

You can configure the following event notification methods:

- **E-mail notification.** Sends events to e-mail addresses. The group attempts to combine multiple events into a single e-mail message to reduce the number of sent messages. A message will be delayed three minutes at most.

See *Configuring E-Mail Event Notification* on page 3-21.

- **E-Mail Home.** E-Mail Home enables Technical Support to rapidly respond to issues and assist customers. If a hardware component (disk, control module, fan, or power supply) fails or if firmware on an array is updated, the group automatically sends e-mail to Technical Support and a local e-mail address you specify.

E-Mail Home is available to all PS Series customers, but response time and assistance depends on the validity and level of your support contract.

See *Configuring E-Mail Home Notification* on page 3-23.

- **Remote server logging.** Logs events to the syslog facility on a remote server that must be configured to accept remote log files. When an event is generated, it is immediately sent to the syslog servers.

See *Configuring Syslog Event Notification* on page 3-24.

- **Informational event display.** Controls whether to display live messages with a priority of INFO in the Group Events window and on the console.

See *Displaying Live Informational Messages* on page 3-24.


For any type of e-mail notification, the group must have access to a Simple Mail Transfer Protocol (SMTP) server or e-mail relay. Likewise, remote server logging requires the availability of the server.

Configuring E-Mail Event Notification

To configure e-mail event notification, follow these steps:

1. Click `Group Configuration > Notifications` tab. The Group Notifications window (Figure 3-12) appears.

Figure 3-12: Group Notifications



2. In the E-Mail Event Notifications panel, select **Send E-mail to addresses** and click **Add**. In the Add List Item dialog box, specify an e-mail address and click **OK**. You can specify up to five addresses. All will receive events.
3. Under **Event priorities**, select the event priorities for which you want to be notified. These priorities will apply to all the specified e-mail addresses.
4. Specify one to three SMTP servers or e-mail relays to handle e-mail forwarding. Under **E-Mail Configuration Settings**, click **Add**. In the Add List Item dialog box, specify an e-mail address and click **OK**. Use the `ip_address:port` format to specify a port number if you are using a port other than the default, 25.

You can specify up to three addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows to change the address order.

To modify or delete an IP address, select the address and click `Modify` or `Delete`, as desired.

5. To ensure that the e-mail recipient does not discard or reject the notifications, under `E-Mail Configuration Settings`, in the `Sender in e-mail address` field, specify the e-mail address that will appear in the “From” field in the notification e-mail.

To test the e-mail notification, you can log out of the GUI and then log in again. Each configured e-mail address should receive the message.

Configuring E-Mail Home Notification

If you configure E-Mail Home and a hardware component (disk, control module, fan, or power supply) fails or firmware on an array is updated, the group automatically sends e-mail to Technical Support and to a local e-mail address that you specify.

To configure E-Mail Home notification, follow these steps:

1. Click `Group Configuration > Notifications` tab. The `Group Notifications` window (Figure 3-12) appears.
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, specify the e-mail address that will receive E-Mail Home notifications (in addition to Technical Support). This is required.
4. Specify one to three SMTP servers or e-mail relays to handle e-mail forwarding. Under `E-Mail Configuration Settings`, click `Add`. In the `Add List Item` dialog box, specify an e-mail address and click `OK`. Use the `ip_address:port` format to specify a port number if you are using a port other than the default, 25.

You can specify up to three addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows to change the address order.

To modify or delete an IP address, select the address and click `Modify` or `Delete`, as desired.

5. To ensure that the e-mail recipient does not discard or reject the notifications, under E-Mail Configuration Settings, in the `Sender in e-mail address` field, specify the e-mail address that will appear in the “From” field in the notification e-mail.

When you first enable E-Mail Home, you will receive a confirmation message stating that the functionality is enabled. If you do not receive this message, check the configuration to ensure that the local e-mail address is correctly configured. Also, check the PS Series event log. If the configuration is correct and no errors are logged, contact your array support provider.

When Technical Support receives an E-Mail Home message, if you have a service agreement, Technical Support will respond to the local e-mail with information on how to proceed.

Configuring Syslog Event Notification

To configure syslog event notification, follow these steps:

1. Click `Group Configuration > Notifications` tab. The Group Notifications window (Figure 3-12) appears.
2. In the Event Logs panel, select `Send events to syslog servers` and click `Add`. In the Add List Item dialog box, specify an IP address for a remote syslog server and click `OK`. You can specify up to three IP addresses, which can include ports. All remote hosts will receive events.

To modify or delete an IP address, select the address and click `Modify` or `Delete`, as desired.

3. Under Event priorities, select the event priorities for which you want to be notified. These priorities will apply to all the specified syslog servers.

Displaying Live Informational Messages

To display live messages with a priority of INFO in the Group Events window and on the console, click:

`Group Configuration > Notifications` tab

The Group Notifications window (Figure 3-12) appears.

Under Event Logs, select `Enable live informational messages`.

To disable the displaying of live informational messages, de-select `Enable live informational messages`.

Note: Informational messages will still be logged to the event message database even if you disable the displaying of live messages.

Using CHAP to Control iSCSI Access to Volumes

Challenge Handshake Authentication Protocol (CHAP) is a network login protocol that uses a challenge-response mechanism. You can use CHAP authentication to restrict iSCSI access to volumes and snapshots to hosts that supply the correct account name and password (or “secret”) combination.

Using CHAP authentication can facilitate the management of access controls because it restricts access through account names and passwords, instead of IP addresses or iSCSI initiator names.

The iSCSI protocol supports two levels of CHAP authentication, initiator and target authentication, described as follows:

- **Initiator authentication.** The iSCSI initiator (host) is authenticated by the iSCSI target (volume or snapshot). When an initiator tries to connect to a target (manually or through discovery), it provides a user name and password to the target. Some implementations refer to the password as a “secret.”

The target checks whether the supplied user name matches an entry in an access control record for the volume. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for information on setting up access control records.

If a match exists, a check is performed to determine if the user name and password combination matches an entry in a CHAP database. If the supplied user name and password match a CHAP database entry, the host will be able to connect to the volume or snapshot.

You can implement a CHAP database in a group in the following ways:

- Local CHAP accounts configured in the group. Local CHAP is not dependent on any external system and is generally used if you have only a few CHAP users. Up to 100 local CHAP accounts are supported. See *Using Local CHAP Accounts for Initiator Authentication* on page 3-26 for more information.

- External RADIUS server whose IP address is known to the group. An external RADIUS server is beneficial if you are managing a large number of CHAP user names and passwords. However, the availability of the server will affect host access. You can configure the group to use multiple RADIUS servers for high availability. See *Using External RADIUS Servers for Initiator Authentication* on page 3-28 for more information.
- **Target authentication.** Each iSCSI target presented by the group is authenticated by the iSCSI initiator. When an initiator tries to connect to a target, the target provides a user name and password to the initiator. The initiator compares the supplied user name and password to information it holds. If they match, based on a hashing algorithm, the initiator can connect to the target.

On the group side, target authentication is always enabled, although you can modify the password and account name as needed. The iSCSI initiator settings determine whether the target authentication is enforced. See *Configuring Target Authentication* on page 3-29 for more information.

Initiator authentication can be implemented without target authentication. However, target authentication can be implemented only if initiator authentication is also implemented. When used together, initiator and target authentication provide **mutual authentication**; that is, both the initiator and the target authenticate each other.

Using Local CHAP Accounts for Initiator Authentication


Local CHAP is not dependent on any external system and is generally used if you have only a few CHAP users. Up to 100 local CHAP accounts are supported.

To add local CHAP accounts for initiator authentication, click:

Group Configuration > iSCSI tab

The Group iSCSI window appears (Figure 3-13).

Figure 3-13: Group iSCSI



In the Local CHAP Accounts panel, click Add. The Add CHAP Account dialog box appears (Figure 3-14). Specify a CHAP user name and password. If you leave the password field blank, a password will be generated automatically.

Figure 3-14: Add CHAP Account



Note: For optimal security, passwords used in CHAP authentication should contain at least 12 characters (preferably random). Individual initiators may have their own rules and restrictions for length and format. Consult the initiator documentation for details.

By default, a CHAP account is enabled. To disable the account, de-select `Enable CHAP account`. Later, you can enable the account to activate it. Click `OK` to create the account.

In addition, to prevent hosts from discovering targets for which they are not authorized, in the iSCSI Discovery panel of the Group iSCSI window (Figure 3-13), select `Prevent unauthorized hosts from discovering targets`. Otherwise, initiators that support discovery will attempt to log in to the target, even if they do not have the right access credentials, resulting in a large number of events logged in the group and an inefficient use of resources.

To modify or delete a local CHAP account, in the Local CHAP Accounts panel, select the account name and click `Modify` or `Delete`, as desired.

After creating the CHAP account, create an access control record for a volume and specify the CHAP user name in the record. To access the volume, a host must supply the user name and its password. A host must meet *all* the requirements in one access control record to access the volume. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.

Using External RADIUS Servers for Initiator Authentication

Before configuring RADIUS authentication servers in the Group Manager GUI, you must set up accounts on the RADIUS servers as described in *Setting Up External RADIUS Authentication Servers* on page 3-13.

To use an external RADIUS server for initiator authentication, follow these steps:

1. Click `Group Configuration > iSCSI` tab. The Group iSCSI window appears (Figure 3-13).
2. In the iSCSI Authentication panel, select `Enable RADIUS authentication for iSCSI initiators`.
3. If you want the group to check for a matching local CHAP account before checking the RADIUS server, select `Consult locally defined CHAP accounts first`.

4. Specify the RADIUS authentication servers. Click `RADIUS settings`. The RADIUS Settings dialog box (Figure 3-9) appears. Follow the instructions in *Setting Up External RADIUS Authentication Servers* on page 3-13.

Note: A RADIUS server must be set up and available on the network in order for the group to access it. Be sure that the RADIUS server is highly available. Downtime will disrupt host access for any volume that requires CHAP authentication.

5. Optionally, to prevent hosts from discovering targets for which they are not authorized, in the iSCSI Discovery panel of the Group iSCSI window, select `Prevent unauthorized hosts from discovering targets`. Otherwise, initiators that support discovery will attempt to log in to the target, even if they do not have the right access credentials, resulting in a large number of events logged in the group and an inefficient use of resources.

After specifying a RADIUS server, create an access control record for a volume and specify a CHAP user name (already configured on the RADIUS server) in the record. To access the volume, a host must supply the user name and its password. A host must meet *all* the requirements in one access control record to access the volume. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.

Configuring Target Authentication

To configure the group for target authentication, in which each iSCSI target presented by the group is authenticated by the iSCSI initiator, you must first set up initiator authentication, as described in *Using Local CHAP Accounts for Initiator Authentication* on page 3-26 or *Using External RADIUS Servers for Initiator Authentication* on page 3-28. Using both initiator and target authentication is called mutual authentication.

Note: Target authentication is always enabled on the group. The iSCSI initiator settings determine whether target authentication is enforced.

To display the target authentication user name and password for the group, click:

Group Configuration > iSCSI tab

The Group iSCSI window appears (Figure 3-13).

In the iSCSI Authentication panel, under Target authentication, are shown the group's default account name and password (referred to by some initiators as a "secret") required

for target authentication. The password displayed here *must* match the information held by the iSCSI initiator.

If you want to change the user name or password, click **Modify**. In the Modify Target CHAP Account dialog box (Figure 3-15), modify the information and click **OK**. Passwords should include at least 12 characters, preferably random.

Figure 3-15: Modify Target CHAP Account



Using iSNS for iSCSI Device Management

You can configure a group to use an iSNS (Internet Storage Name Service) server, which facilitates the automated discovery, management, and configuration of iSCSI devices and other devices on a TCP/IP network.

An iSNS server provides a central management point for the group by dynamically maintaining up-to-date information about the iSCSI target names for group volumes. Once the IP address of an iSNS server is entered in an iSCSI initiator's configuration utility, the setting is persistent across initiator sessions. There is no need to enter the group IP address or iSCSI target name in the utility.

Note: An iSNS server must already be set up and available on the network in order for the group to access it.

To configure iSNS servers in the group, click:

Group Configuration > iSCSI tab

The Group iSCSI window appears (Figure 3-13).

In the iSCSI Discovery panel, click **Add**. In the Add List Item dialog box, specify the IP address for an iSNS server. You can specify up to three IP addresses. Only one server is used at one time. The first server specified is the default server. The other servers are used, in the order specified, if the default server is not available. Use the up and down arrows to change the address order.

Use the *ip_address:port* format to specify a port number if you are using a port other than the default, 3205.

To modify or delete an address, select the address and click **Modify** or **Delete**, as needed.

Using SNMP to Monitor a Group

To monitor a group using Simple Network Management Protocol (SNMP), you can configure SNMP read-only access to a group. See *Configuring SNMP Read-Only Access* on page 3-31.

You can also specify IP addresses to receive SNMP traps. See *Sending SNMP Traps to a Destination* on page 3-32.


Configuring SNMP Read-Only Access

To configure SNMP read-only access, click:

Group Configuration > SNMP tab

The Group SNMP window appears (Figure 3-16).

Figure 3-16: Group SNMP



To allow access from an SNMP management console to the group, you must specify at least one read-only SNMP community name, which acts as a password.

To specify an SNMP community name, in the SNMP Access panel click **Add**. In the Add list item dialog box, specify a unique SNMP read-only community name, up to 64 alphanumeric characters, and click **OK**. You can specify up to five read-only community names for the SNMP agent to use to authenticate SNMP requests.

To modify or delete a community name, select the name and click **Modify** or **Delete**, as desired.

Sending SNMP Traps to a Destination

Traps are unsolicited event messages provided to the management console by the agent. The PS Series array SNMP traps can be categorized into equipment issues and security issues.

To send SNMP traps to a destination, click:

Group Configuration > SNMP tab

The Group SNMP window appears (Figure 3-16).

First, configure an SNMP community name, as described in *Configuring SNMP Read-Only Access* on page 3-31.

Then, in the SNMP Traps panel, click **Add**. In the Add list item dialog box, specify the IP address where SNMP traps (event messages) will be sent and click **OK**. You can specify up to five IP addresses. All will receive traps.

To modify or delete a trap destination, select the IP address and click **Modify** or **Delete**, as desired.

The SNMP trap community name is used when sending SNMP traps. The default name is `SNMP-trap`. To change the SNMP trap community name, in the Group SNMP window (Figure 3-16) enter a new name in the `SNMP Trap community name` field. Names must be unique and contain up to 64 alphanumeric characters. Commas are not allowed.

The PS Series array MIBs (Management Information Bases) contain information about the traps and trap thresholds. If you have a support account, you can download and view the MIBs from the Technical Support website. Log in to a support account and click on **Downloads**. Click the link for your PS Series Firmware version, then click **MIBS**.

The MIBs are collected in an archive. You can open this archive and view the individual MIBs or you can download the archive and store it on your system. To use MIBs, you must install them on a management station, as required by the management tools.

Not all traps apply to all array models.

PS Series array traps include the following:

- `linkUp`
- `linkDown`
- `authenticationFailure`
- `coldStart`
- `warmStart`
- `iscsiTgtLoginFailure`

- iscsiIntrLoginFailure
- iscsiInstSessionFailure
- eqlMemberHealthTempSensorHighThreshold
- eqlMemberHealthTempSensorLowThreshold
- eqlMemberHealthFanSpeedHighThreshold
- eqlMemberHealthFanSpeedLowThreshold
- eqlMemberHealthPowerSupplyFanFailure
- eqlMemberHealthPowerSupplyFailure
- eqlMemberHealthRAIDSetDoubleFaulted
- eqlMemberHealthBothFanTraysRemoved
- eqlMemberHealthRAIDlostCache
- eqlMemberHealthFanTrayRemoved
- eqlMemberHealthRAIDSetLostBlkTableFull
- eqlMemberHealthBatteryLessThan72Hours
- eqlMemberHealthRaidOrphanCache
- eqlMemberHealthRaidMultipleRaidSets
- eqlMemberHealthNVRAMBatteryFailed
- eqlMemberHealthhwComponentFailedCrit
- eqlMemberHealthincompatControlModule
- eqlMemberHealthlowAmbientTemp
- eqlMemberHealthopsPanelFailure
- eqlMemberHealthemmLinkFailure
- eqlMemberHealthhighBatteryTemperature
- eqlDiskStatusChange
- scsiTgtDevicesStatusChanged
- scsiLuStatusChanged

Modifying Access Controls for Windows Services


Microsoft Windows systems that are running Microsoft services, such as Virtual Disk Service (VDS) and Volume Shadow Copy Service (VSS), must be able to access a PS Series group for management operations.

Note: Access controls set up in the group must match the access controls on the Windows system, otherwise Microsoft service access will fail.

To use the Group Manager GUI to modify an access control record for Microsoft service access to the group, click `Group Configuration > VDS/VSS tab`.

The Group VDS/VSS window appears (Figure 3-17).

Figure 3-17: Group VDS/VSS



Select a record to display details.

To add a record, click `Add`.

In the Add Access Control Record dialog box (Figure 3-18), specify an IP address, iSCSI initiator name, or CHAP user name (or any combination). A Windows system must meet all the requirements specified in an access control record in order for Microsoft services to access the group. For example, the CHAP user name stored on the server must match the CHAP user name in the record.

Figure 3-18: Add Access Control Record



After entering the information, click **OK**. The access control record will appear in the Group VDS/VSS window.

Note: If you are using CHAP for authentication, be sure to set up a matching local CHAP account, as described in *Using CHAP to Control iSCSI Access to Volumes* on page 3-25.

To modify or delete an existing access control record, select the record and click **Modify** or **Delete**.

Modifying the Group-Wide Volume Settings

When you create a volume, the group-wide defaults for snapshot space, thin-provisioning, and iSCSI aliases are applied unless you explicitly override them. As needed, you can change the group-wide defaults.


Note: Changing the group-wide, default volume settings will affect only new volumes.

To modify the group-wide, default volume settings, click:

Group Configuration > Defaults tab

The Group Defaults window appears (Figure 3-19).

Figure 3-19: Group Defaults



In the Group Defaults window, you can modify the following default settings:

- **Snapshot reserve.** If you want to create snapshots of volumes, specify the amount of pool space, as a percentage of the volume reserve, to reserve for snapshots. The default is 100% of the volume reserve. If you enter a value, this space is reserved even if you never create a snapshot. If no snapshot space is reserved for a volume, snapshot creation will fail.
- **Snapshot space warning limit.** In the Warn when in-use space reaches field, specify the percentage of the snapshot reserve that, when consumed, results in an alarm. The default is 90%.

For example, if a volume size is 200 MB and you reserve 100% of the volume size for snapshots, using the default limit of 90%, a warning will be generated when the amount of in-use snapshot reserve reaches 180 MB.

- **Snapshot space recovery policy.** Select the action to take when the space reserved for snapshots has been exceeded: either delete the oldest snapshots until sufficient space

is recovered (the default) or set the volume and snapshots offline. Active iSCSI connections will be terminated before a snapshot is automatically deleted.

If a snapshot deletion is imminent, you can avoid losing its data by cloning the snapshot.

- **In-use space warning limit.** For thin-provisioned volumes, you can change the default in-use space warning limit. When the in-use volume reserve reaches this percentage of the volume size, a warning event message is generated. The default is 60% of the volume size.

For example, if you create a thin-provisioned volume with a size of 500 GB and set the warning limit to 75%, a warning is generated when the amount of used space in the volume reaches 75% of 500 GB, or 375 GB.

- **Maximum in-user space limit.** For thin-provisioned volumes, you can change the default maximum in-use space limit, which is the maximum percentage of the volume size that can be in use. The volume reserve cannot increase above this value. When in-use space reaches the maximum, the volume is set offline. The default is 100% of the volume size.

For example, if you create a thin-provisioned volume with a size of 500 GB and set the maximum in-use space limit to 90%, the volume is set offline when the amount of used space in the volume reaches 90% of 500 GB, or 450 GB.

- **iSCSI target alias.** If you select `Use volume name for default iSCSI target alias`, the iSCSI target alias (public name) for a volume or snapshot will be the same as the volume or snapshot name. If you de-select the checkbox, no alias is automatically assigned. The alias is seen on the network by some iSCSI initiators. It can be used by administrators to help identify the iSCSI target.

Controlling Performance Load Balancing

A PS Series group provides several types of automatic load balancing across storage pool resources, including capacity, network connection, and performance load balancing.

When performance load balancing is enabled, the group attempts to place volume data on a pool member with a RAID configuration that is optimal for volume performance. If you disable performance load balancing, volume data is balanced across the pool members only by capacity.


Note: Manually assigning a RAID preference to a volume or binding a volume to a member will disable automatic performance load balancing for that volume. See *Setting a Volume RAID Preference* on page 5-35 and *Binding a Volume to a Specific Member* on page 5-37 for more information.

To enable or disable performance load balancing, click:

Group Configuration > Advanced tab

The Group Advanced window appears (Figure 3-20).

Figure 3-20: Group Advanced



In the Advanced Features panel, select **Enable performance load balancing in pools** to enable automatic load balancing. De-select the checkbox to disable it.

Preventing Member Firmware Downgrades

In some cases, you might want to prevent a PS series group and its members from running an older version of firmware. Specifically, if you have a group in which one or more members were upgraded to Version 3.2 (or higher) and you want to use the new features available in the latest release, you must prevent the group members from going back to a previous firmware version in which those features are not supported.

To prevent the group members from returning to a previous firmware version, click:

Group Configuration > Advanced tab

The Group Advanced window appears (Figure 3-20).

If the group already disallows downgrades to earlier firmware versions, the `Disallow downgrades` option will be grayed out and the GUI will display the message `Firmware downgrades are disallowed`.

If not, click `Disallow downgrades`. This operation cannot be reversed.

4 Storage Pool Management

Dividing the overall PS Series group space into multiple pools provides a greater degree of control over storage resources and can be useful in environments where a variety of applications compete for SAN storage. Using pools is optional and requires a multi-member PS Series group.

Only group administrators can create pools and move objects across pools. Only group or pool administrators can modify a pool's properties. For more information, see *Managing Group Administration Accounts* on page 3-9.

Tasks related to storage pool management include:

- *Using Storage Pools to Separate SAN Storage* on page 4-1
- *Planning Storage Pools* on page 4-4
- *Creating a Storage Pool* on page 4-5
- *Displaying Storage Pool Configuration and Status* on page 4-6
- *Modifying a Storage Pool Name or Description* on page 4-10
- *Merging Storage Pools* on page 4-10
- *Deleting a Storage Pool* on page 4-11

For a detailed assessment of using storage pools, see the related Technical Report, which you can download from the Technical Support website.

Using Storage Pools to Separate SAN Storage

Storage consolidation through a SAN allows many applications and workloads to use a single pool of storage. Well-designed SANs can improve storage utilization, increase reliability, and dramatically simplify storage management.

While storage consolidation makes it easier to manage, provision, and back up data, it can limit your ability to meet other business objectives. For example, environments with conflicting or competing storage requirements may need more flexibility in allocating SAN resources—while still maintaining a single-system view of the storage.


To meet the need for flexible and efficient storage, you can divide PS Series group space into multiple storage pools. Using “SAN within a SAN” technology, administrators can easily segregate workloads, while retaining the advantages of storage consolidation.

For example, you can separate storage by application, service level, disk type, cost, location, or department. You can assign mission critical applications to their own storage hardware resources—ensuring fast and consistent performance—and assign disaster recovery replicas, backups, and archived data to different resources.

In order to separate storage space in a PS Series group, you need multiple members. You can create up to four storage pools in a group and assign each member to a pool.


If a pool has sufficient space, you can assign volumes to the pool. Only the resources provided by the pool members are available to the volumes.

For example, Figure 4-1 shows a single-member group with one pool (the default pool). Volume data is automatically balanced across the member disks.




When you add a member to a group, you can assign the new member to the default pool, or you can create another pool and assign the member to the new pool. Volumes can then be assigned to the new pool.

For example, Figure 4-2 shows a three-member, single-pool group. Volume data is automatically balanced across the pool members. In addition, performance load balancing will attempt to place volume data on pool members with the RAID level that is optimal for the application, as determined by internal group metrics.

Figure 4-2: Three-Member, Single-Pool Group

To separate storage in a multi-member group, you must create additional pools and assign members and volumes to the pools. Only the storage space and network resources provided by the pool members are available to the volumes assigned to the pool. As usual, automatic performance and capacity balancing takes place—but only among the pool members.

Figure 4-3 shows a three-member, two-pool group.

Figure 4-3: Three Member, Two-Pool Group

Members and volumes can be assigned only to one pool at a time. As capacity needs change, you can move members and volumes from one pool to another. When a member is moved to a different pool, all volume data on the member must be moved to the remaining pool members. Therefore, the remaining pool members must have enough space to store the volume data.

When moving members and volumes between pools, data movement occurs online, enabling administrators to easily optimize business objectives without downtime or after-hours maintenance, thus improving storage utilization and service levels.

Planning Storage Pools

When planning your storage pool configuration, you will want to place volumes in the pool that provides the appropriate capacity and level of performance and service. You may need to prioritize applications, determine how you want to segregate them, and ensure you have sufficient capacity in each pool.

To create a successful storage configuration, identify the following:

- **Capacity and performance of each member.** You should consider the following factors:
 - Disk type, either Serial Attached SCSI (SAS) or Serial Advanced Technology Attachment (SATA)
 - Disk size
 - Disk speed
 - Member RAID level

Within a pool, it is recommended that 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, both members should be configured with the same RAID level.

- **Disk space and performance needs of each application.** You should determine which applications deserve priority and identify their disk space and network bandwidth needs and performance characteristics. For example, some applications perform many random data transfers, while others perform a small number of large sequential data transfers. In addition, you should determine which volumes do not require the highest performance, such as rarely-accessed volumes that are used only to archive data.
- **How the pools will be organized.** Identify the number of pools and which members will be assigned to each pool, based on the member performance characteristics and capacities and the requirements of the volumes. For example, you may want to organize your pools according to disk speed. You could also organize pools according to member RAID level.
- **Application volumes that will be assigned to each pool.** Volumes with the highest service level needs should be assigned to a pool that meets those needs. The volumes used by an application can be assigned to different pools, as needed.

When assigning volumes to pools, the RAID level of the pool members may be important. For a workload characterized by many small random transfers, RAID 10 may provide the highest throughput, while a workload characterized by many large

sequential transfers may perform equally well with either RAID 10 or RAID 50. If you do not know the optimal RAID level for a volume, be sure that the pool to which the volume is assigned contains members with a variety of RAID levels, so automatic performance load balancing can occur.

You should consider future demands for capacity and performance when planning the storage pool configuration.

In addition, monitor the group and applications to ensure that the storage configuration provides adequate service levels across various workloads.


As needed, you can move members and volumes from one pool to another—with no impact on availability or performance. For example, if the capacity of a pool is not sufficient, you can either add a new member to the pool or move an existing member to the pool.

For more information on using pools, see the related Technical Report, available on the Technical Support website.

Creating a Storage Pool

To create a storage pool, in the Activities panel of the Group Summary window (Figure 3-1), click `Create storage pool`. The Create Storage Pool dialog box (Figure 4-4) appears.

Figure 4-4: Create Storage Pool



Specify a pool name and optional description and click OK.


To assign a member to a storage pool, see *Moving a Member Across Pools* on page 9-27.

To assign a volume to a storage pool, see *Moving a Volume Across Pools* on page 5-29.

Displaying Storage Pool Configuration and Status

To display the storage pools in a group, in the far left panel, click Storage Pool. The Storage Pool Summary window (Figure 4-5) appears.

Figure 4-5: Storage Pool Summary



In the Group Disk Space panel, select the type of information you want to display. For example, you can select the following:

- Storage pool capacity – Percentage of group space in each pool and the capacity of each pool.
- Free storage pool space – Free space in each pool.

Note: If a pool's free space is low, you can increase free space by moving volumes from the low-space pool to a different pool, as described in *Moving a Volume Across Pools* on page 5-29. You can expand pool capacity by moving an


existing member to the low-space pool, as described in *Moving a Member Across Pools* on page 9-27, or adding a new member to the group and assigning it to the low-space pool, as described in *Adding Members to an Existing Group* on page 2-5.

- Group space utilization – Group space allocated to volumes, snapshots, local replication reserve, and delegated space and the percentage of free group space.
- RAID level space distribution – Group space configured with each RAID level (RAID 10, RAID 50, or RAID 5).

For each pool, the Storage Pools panel shows total pool capacity, free space, pool space configured for each RAID level, and the number of members, volumes, and snapshots. Select a pool to display related tasks in the Activities panel.

To display details about a pool, select the pool name in the far left panel or double-click the pool in the Storage Pools panel. The Storage Pool Status window appears (Figure 4-6).

Figure 4-6: Storage Pool Status



The General Pool Information window displays the status of the pool members and the number of volumes and snapshots in the pool, in addition to the number of volumes and snapshots that are online and in use (that is, have a current iSCSI connection).

In the Pool Disk Space panel, select the type of information you want to display:

- Pool space utilization – Pool space allocated to volumes, snapshots, local replication reserve, and delegated space and the percentage of free pool space.
- RAID level space distribution – Pool space configured with each RAID level (RAID 10, RAID 5, and RAID 50).
- Snapshot space utilization – Used and free pool space reserved for volume snapshots.
- Delegated space utilization – Used and free pool space delegated to all replication partners.


Note: If free delegated space is low, monitor the replication partners, as described in *Displaying Replication Activity Between Partners* on page 7-32, to determine which partner needs more delegated space. See *Modifying a Replication Partner* on page 7-17 for information on increasing delegated space.

The Pool Members panel lists the members that comprise the pool. For each member, the panel shows the member's name, capacity, free space, RAID level, number of disks and their type (SAS or SATA), status, firmware version, and iSCSI connections. Select a member to display related tasks in the Activities panel.

Note: If a pool member is offline, be sure to investigate. See *Displaying Member Configuration and Status* on page 9-6 for information about monitoring members.

To display detailed information about the volumes that have been assigned to the pool, click the `Volumes` tab. The Storage Pool Volumes window appears (Figure 4-7).

Figure 4-7: Storage Pool Volumes



For each volume, the Pool Volumes panel displays the volume name, reported size (seen by iSCSI initiators), volume reserve (allocated space), snapshot reserve, status, replication partner, number of snapshots, and iSCSI connections. The volume icon indicates whether thin provisioning is enabled on the volume.

Select a volume to display related tasks in the Activities panel.

Notes: A volume can be offline due to a failure or administrator action, or, if thin provisioning is enabled, because the volume has reached its maximum growth limit. If a volume is offline, you should investigate. See *Displaying Volume Status and Configuration* on page 5-11 for information about monitoring volumes.


If a volume's free snapshot reserve is low, consider increasing the snapshot reserve. If the snapshot reserve is exceeded, the snapshot reserve recovery policy will be enforced and either snapshots will be deleted or the volume will be set offline. See *Modifying Snapshot Settings* on page 6-3 for information about increasing snapshot reserve.

Modifying a Storage Pool Name or Description

You can change the name or description of a storage pool. If you change the name of the default pool, the renamed pool will still behave as the default pool (that is, members and volumes will be assigned to this pool by default).

To change the name or description of a storage pool, in the far left panel, expand `Storage pools` and select the pool. In the Activities panel, click `Modify pool settings`. The `Modify Pool Settings` dialog box appears (Figure 4-8).

Figure 4-8: Modify Pool Settings



Modify the pool name and description, as needed, and click `OK`.

Merging Storage Pools

You can merge any storage pool except the default pool into any other existing pool.


Merging a storage pool moves the pool members and all their volumes and snapshots into the destination pool. The empty pool is then deleted.

Note: You cannot create a new pool by merging pools.

To merge a pool, in the far left panel, expand `Storage pools` and select the pool you want to merge into another pool. In the Activities panel, click `Merge storage pool`.

The Merge Storage Pools dialog box appears (Figure 4-9).

Figure 4-9: Merge Storage Pools



Select the destination pool and click `OK`.

Deleting a Storage Pool

If you delete a storage pool, all the members and volumes in the pool will be automatically and instantly moved to the default pool. If desired, you can then assign the members and volumes to a different pool, if one is available.

Note: You cannot delete the default pool.

To delete a storage pool, in the far left panel, expand `Storage pools` and select the pool. In the Activities panel, click `Delete storage pool`. Confirm that you want to delete the pool.

5 Volume Management

To access storage in a PS Series group, you allocate portions of a storage pool to volumes. Hosts with iSCSI initiators can then connect to the volumes.

Only group administrators or pool administrators can create or modify a volume. For more information, see *Managing Group Administration Accounts* on page 3-9.

Volume administration topics include the following:

- *Introduction to Volumes* on page 5-1
- *Creating Volumes* on page 5-3
- *Understanding Thin Provisioning* on page 5-8
- *Displaying Volume Status and Configuration* on page 5-11
- *Managing Volume Collections* on page 5-15
- *Managing Access Controls for Volumes and Snapshots* on page 5-23
- *Connecting Hosts to Volumes and Snapshots* on page 5-27
- *Modifying Volume Settings* on page 5-29
- *Cloning Volumes* on page 5-38
- *Managing Schedules for Snapshots and Replicas* on page 5-41
- *Setting a Volume Offline or Online* on page 5-48
- *Setting Volume Permission* on page 5-48
- *Deleting Volumes* on page 5-49
- *Handling Lost Blocks* on page 5-49

See Chapter 6, *Snapshot Management* and Chapter 7, *Basic Replication* for information about managing volume snapshots and replicas.

Introduction to Volumes

In a PS Series group, each volume is assigned a size, storage pool, and access controls. Optionally, you can reserve snapshot or replication space for a volume. Volume data is automatically load balanced across pool members, as needed.

Volumes are seen on the network as **iSCSI targets**. Only hosts with an **iSCSI initiator** and the correct access credentials can access a volume. The group automatically generates an iSCSI target name for each volume. This is the name that iSCSI initiators use to access the volume.

An iSCSI target name consists of a prefix (`iqn.2001-05.com.equallogic`), followed by a generated string, followed by the volume name. An iSCSI target name must always adhere to the iSCSI protocol specification.

An example of an iSCSI target name for a volume named `db3` is shown next:

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

For backup or disaster recovery purposes, you can create volume **snapshots** or **replicate** a volume to a different group. You can also **clone** a volume, snapshot, or replica to create an exact copy.

Related volumes can be grouped into a **volume collection**. You can then create snapshots or replicas of the volumes in one operation. In addition, schedules can be set up to create snapshots or replicas of volumes and volume collections at a specific time or on a regular basis.

Optionally, you can use **thin provisioning** to improve storage resource utilization by logically over-allocating SAN volumes. Thin provisioning enables you to satisfy current and future storage needs and eliminates the difficulty of adding storage to operating systems or applications that do not support online expansion.

A volume has two sizes associated with it. The **reported size** is the size that is seen by iSCSI initiators. Systems connected to the volume will associate this amount of space with the volume. The **volume reserve** is the amount of space that is allocated to the volume. For a volume without thin provisioning, the reported size and the volume reserve will be the same. For a thin-provisioned volume, the reported size will be greater than (or equal to) the volume reserve, because the volume is not fully allocated.

See *Understanding Thin Provisioning* on page 5-8 for more information about thin provisioning. See Chapter 6, *Snapshot Management* and Chapter 7, *Basic Replication* for more information about snapshots and replicas.

Creating Volumes

When you create a volume, you specify the volume size (reported size that is seen by initiators), snapshot reserve, access control information, and other attributes. An iSCSI target name will be automatically created for the volume.

To create a volume, click:

Volumes > Create volume

The Create Volume - General Settings dialog box appears (Figure 5-1).

Figure 5-1: Create Volume - General Settings

Create volume

Step 1 - Volume Settings

General properties:

* Volume name:

Description:

Storage pool assignment:

Storage Pool	Capacity	Free	Drives
☒ default	5.23 TB	3.48 TB	SATA
☐ usermail	7.85 TB	1.22 TB	SAS

Navigation: Back Next Finish Cancel

Java Application Window


You can specify information in the following fields:

- **Volume name** – Unique name, up to 64 alphanumeric characters (including periods, hyphens, and colons), used to identify the volume for administrative purposes. The volume name appears at the end of the iSCSI target name that is generated for the volume. Host access to the volume is always through the iSCSI target name, not the volume name.
- **Description** – The volume description is optional.

- **Storage pool** – All volume data will be restricted to the members that make up the pool. By default—or if only one pool exists—the volume will be assigned to the default pool. If multiple pools exist, you can assign the volume to a different pool. See Chapter 4, *Storage Pool Management* for information about setting up pools.

Click **Next** to display the Create Volume – Space Reserve dialog box (Figure 5-4).

Figure 5-2: Create Volume - Space Reserve



Create volume

Step 2 - Space reserve

Volume space:

* Volume size: 25 ☐ MB ☒ GB (max. 3.48 TB)

☐ Thin provisioned volume (use sliders below to adjust settings)

Snapshot space:

* Snapshot reserve: 100 % of volume reserve

Reported volume size 25.0 GB

Free 25.0 GB

Estimated changes in storage pool default

	Current	New	Change
Volume reserve	137.9 GB	162.91 GB	25.0 GB
Snapshot reserve	137.9 GB	162.91 GB	25.0 GB
Replication reserve	137.9 GB	137.9 GB	0 MB
Delegated space	1.32 TB	1.32 TB	0 MB
Free pool space	3.48 TB	3.43 TB	-50.01 GB

Back Next Finish Cancel

Java Application Window

Specify information in the following fields:

- **Volume size** – This is the reported size of the volume as seen by iSCSI initiators. Be sure to select the correct unit (by default, gigabytes). The minimum volume size is 15 MB. In addition, volume sizes are rounded up to the next multiple of 15.

The largest volume size you can specify appears next to the **Volume size** field. This value depends on the volume's thin provisioning setting.

Note: For a standard volume (without thin provisioning), the volume reserve (allocated space) is the same as the reported size. For a thin-provisioned volume, the volume reserve will be less than or equal to the reported size.

- **Thin provisioned volume** – Select this checkbox to enable thin provisioning on the volume. When selected, slider bars appear in the Reported volume size panel (Figure 5-3). Use them to modify the following default thin-provisioning values:
 - **Minimum volume reserve** – This is the minimum amount of space to allocate to the volume. As the volume is used, more space will be allocated to the volume, and the volume reserve will increase. The default is 10% of the volume size.
 - **In-use space warning value** – When in-use space reaches this value, as a percentage of the volume size, a warning event message is generated. The warning informs you that volume space is being used, enabling you to make adjustments, as needed. The default is the group-wide volume setting, as described in *Modifying the Group-Wide Volume Settings* on page 3-36.
 - **Maximum in-use space value** – When the in-use space reaches this value, as a percentage of the volume size, the volume is set offline. The default is the group-wide volume setting.

See *Understanding Thin Provisioning* on page 5-8 for more information.

Figure 5-3: Volume Create - Space Reserve - Thin-Provisioned Volume

Create volume

Step 2 - Space reserve

Volume space:

* Volume size: 25 MB GB (max. 15.0 TB)

☒ Thin provisioned volume (use sliders below to adjust settings)

Snapshot space:

* Snapshot reserve: 100 % of volume reserve

Reported volume size 25.0 GB

Free: 10% Unreserved 90%

Estimated changes in storage pool default

	Current	New	Change
Volume reserve	137.9 GB	140.41 GB	2.5 GB
Snapshot reserve	137.9 GB	140.41 GB	2.5 GB
Replication reserve	137.9 GB	137.9 GB	0 MB
Delegated space	1.32 TB	1.32 TB	0 MB
Free pool space	3.48 TB	3.48 TB	-5.01 GB

Back Next Finish Cancel

Java Application Window

- **Snapshot reserve** – If you want to create snapshots of the volume, specify the amount of pool space to reserve for snapshots, as a percentage of the volume reserve.

The default is the group-wide volume setting. If you enter a value, this space is reserved even if you never create a snapshot.

To change the snapshot space recovery policy or warning threshold (the defaults are the group-wide volume settings), you must modify the volume's snapshot settings after creating the volume, as described in *Modifying Snapshot Settings* on page 6-3.


The values in the storage pool space tables shown in Figure 5-2 and Figure 5-3 change to show how creating the volume will affect the pool space utilization, as described in Table 5-1.

Table 5-1: Storage Pool Space Table Values for Volumes

Space	Description	New Value
Volume reserve (same as the reported size for volumes without thin provisioning enabled)	Current amount of pool space allocated to volumes	New amount of pool space allocated to volumes after creating the new volume
Snapshot reserve	Current amount of pool space allocated to snapshots of all volumes	New amount of pool space allocated to all snapshots after adding new volume
Replication reserve	Current amount of pool space allocated to local replication reserve for all volumes	New amount of pool space allocated to local replication reserve after adding new volume. No change for a new volume.
Delegated space	Pool space delegated to replication partners	New amount of pool space delegated to replication partners after adding new volume. No change for a new volume.
Free pool space	Space in the pool not used or reserved for any purpose	New amount of free, unused, unreserved pool space after adding new volume

If creating the volume will exceed the capacity of the pool, the table cell showing free pool space will turn red.

Click **Next** to display the Create Volume – iSCSI Access Policy dialog box (Figure 5-4), which enables you to create an access control record for the volume.

Figure 5-4: Create Volume - iSCSI Access Policy

Access control records are used to restrict volume access to hosts that supply a recognized CHAP user name and password (or “secret”) or that match an IP address or iSCSI initiator name (or any combination of the three).

To create an access control record that will apply to the volume and its snapshots, in the Create Volume – iSCSI Access Policy dialog box, click **Restricted access** and one or more of the following:


- **Authenticate using CHAP user name** – Restricts access to hosts that supply the specified CHAP user name (and its associated password or “secret”). The user name must match a local CHAP account or an account on an external RADIUS server, as described in *Using CHAP to Control iSCSI Access to Volumes* on page 3-25.
- **Limit access by IP address** – Restricts access to hosts with the specified initiator IP address (for example, 12.16.22.123). Use asterisks for “wildcards,” if desired (for example, 12.16.*.*). An asterisk can replace an entire octet, but not a digit within an octet.
- **Limit access to iSCSI initiator name** – Restricts access to hosts with the specified iSCSI initiator.

Note: If using IP addresses or iSCSI initiator names to restrict access, be sure to create an access control record for each IP address or iSCSI initiator name presented by an authorized host.

If you do not want to create an access control record at this time, select `No access`. No host will be allowed access to the volume until you create a record. You can set up access control records or modify records after creating the volume, as described in *Managing Access Controls for Volumes and Snapshots* on page 5-23.

After specifying the access information, click `Next` to display the Create Volume - Summary dialog box (Figure 5-5).

Figure 5-5: Create Volume - Summary



If the configuration is correct, click `Finish` to create the volume. To make changes, click `Back`.

Understanding Thin Provisioning

For some PS Series SAN environments, you can use thin provisioning technology to more efficiently provision storage, while still meeting application and user storage needs. With thin provisioning, businesses no longer have to over-buy storage while large amounts of space remain unused.

Normally, storage administrators will fully provision a volume. That is, space equal to the volume size will be allocated from free space. Because not all operating systems support expanding file systems or raw disks, an administrator may need to allocate more storage space than is currently required in order to accommodate possible future growth.

For example, an administrator may create a 100 GB volume for an organization that will use only a small fraction of the volume. Because the full 100 GB has been allocated to the volume, any unused volume space cannot be used for other purposes. If the space is never used, vital storage resources are wasted.

In contrast, a thin-provisioned volume is initially allocated only a portion of the volume size. As data is written to the volume, more space is automatically allocated (if available), and the volume reserve increases up to a user-defined limit. Regular event messages are generated as space is used, giving the administrator the opportunity to make adjustments, as needed.

A volume reserve for a thin-provisioned volume can automatically grow only if free pool space is available. If space is not available, the auto-grow operation will fail. If in-use space consumes all the volume reserve, the volume will be set offline.

Because thin-provisioned volumes are allocated only the storage space that is required, you can more efficiently utilize storage resources, while eliminating the need to perform difficult resize operations on the host.

Note: Thin provisioning is not always appropriate or desirable in an IT environment. It is most effectively used when you know how a volume will grow over time, the growth is predictable, and users do not need immediate, guaranteed access to the full volume size. Administrators must monitor thin-provisioned volumes and must be prepared to increase storage space by adding or moving arrays to the pool where the volume resides.

If you must guarantee the full volume size, thin provisioning is not recommended for the volume.

In a PS Series group, there are four user-settable parameters that control space utilization in a thin-provisioned volume:

- Volume size – Size that the group reports to iSCSI initiators (shown as reported size in volume status). Systems connected to the volume will associate this amount of space with the volume.
- Minimum volume reserve – Minimum amount of pool space that the group will allocate to the volume. The default is 10% of the volume size.

As data is written to a thin-provisioned volume, volume reserve is consumed. When free volume reserve is less than 6 GB, the group will attempt to allocate additional pool space to the volume reserve, which can increase in size up to the maximum in-use space value described below. Space is allocated in increments of 10 GB, depending on various factors, including the amount of free pool space.

If there is not enough free pool space to increase the volume reserve, the volume reserve will not increase. If in-use space consumes all the volume reserve, the volume will be set offline. If free pool space becomes available, allowing the volume reserve to increase automatically, the volume will be automatically set online.

Note: Snapshot reserve, local replication reserve, and replica reserve for a volume are all based on the current volume reserve, not the reported volume size.

- In-use space warning value – When in-use space reaches this percentage of the volume size, a warning event message is generated. The default is the group-wide volume setting, as described in *Modifying the Group-Wide Volume Settings* on page 3-36.

After the in-use warning value is reached, additional warning messages are generated according to the following policy:

- For volumes larger than 200 GB, additional warnings are generated when the in-use volume reserve increases by every additional 10 GB.
- For volumes smaller than 200 GB, additional warnings are generated when the in-use volume reserve increases by every additional 5%.
- Maximum in-use space – Maximum percentage of the volume size that can be in use. The volume reserve cannot increase above this value. The default is the group-wide volume setting.


If the in-use space reaches the maximum allowed, the volume will be set offline. If pool space is available, and you increase the volume size or the maximum in-use space setting, the volume will be automatically set online.

For information about creating a thin-provisioned volume, see *Creating Volumes* on page 5-3. For information about enabling and disabling thin provisioning on an existing volume, see *Enabling and Disabling Thin Provisioning* on page 5-34.

Displaying Volume Status and Configuration

To display all the volumes in a group and their status, including current iSCSI connections, click **Volumes** in the far left column. The Volume Summary window appears (Figure 5-6).

Figure 5-6: Volume Summary



Volumes are represented by icons that indicate specific attributes, as follows:

- Standard (non-thin-provisioned) volume — A gold-colored cylinder (🔱).
- Thin-provisioned volume — A two-tone, green and gold cylinder (🌿).
- Recovery volume — A gold cylinder with a yellow star on the upper corner and a light bulb on the lower corner, to indicate user information (💡). See *Recovering Data on the Secondary Group* on page 8-28 for information about recovery volumes.

In the Group Disk Space panel, select the type of information you want to display:

- **Group space utilization** — Group space allocated to volumes and snapshots, local replication reserve, delegated partner space for replication, and the amount of free group space.
- **RAID level space distribution** — Group space configured with each RAID level (RAID 10, RAID 5, and RAID 50).
- **Storage pool capacity** — Percentage of group space in each storage pool and the capacity of each pool.
- **Free storage pool space** — Free space in each storage pool.

The Volumes panel provides information about each volume, including the storage pool, reported size, volume reserve, snapshot reserve, number of snapshots, status, replication partner, and number of iSCSI connections. Select a volume name to display volume-related tasks in the Activities panel.

Place the mouse over a volume to display a pop-up message showing both the requested volume status and the current volume status. The requested status for a volume is set by an administrator and described in Table 5-2. The current volume status, also shown in the status column in the Volume Summary window, is described in Table 5-3. If the requested status and the current status are different, you should investigate.

Table 5-2: Requested Volume Status

Status	Description
online	Administrator set the volume online. Authorized hosts can access the volume.
offline	Administrator set the volume offline. Hosts cannot access the volume.
online-lost-blocks	Administrator chose to set the volume online despite lost blocks. Authorized hosts can access the volume. If an application attempts to read a lost block, an error will occur. If the block is re-written, no error will occur, and the block will no longer be marked lost.

Table 5-3: Current Volume Status


Status	Description
online	Volume was set online by an administrator, and no failures have occurred. Authorized hosts can access the volume. Online volumes appear in the far left panel in black text.
offline	Volume was set offline by an administrator, but no failures have occurred. Hosts cannot access the volume. Offline volumes appear in the far left panel in gray text.
offline-snap-reserve-met	Volume was automatically set offline because the reserved snapshot space was exceeded and the snapshot recovery policy dictates that the volume be set offline. Note: To correct this condition, increase the amount of reserved snapshot space, as described in <i>Modifying Snapshot Settings</i> on page 6-3.
offline-thin-max-grow-met	Thin-provisioned volume was automatically set offline because in-use space reached the maximum allowed. Note: To correct this condition, increase the value of the maximum in-use space setting for the volume, as described in <i>Modifying Volume Settings</i> on page 5-29.
offline-nospace-auto-grow	Volume reserve was consumed, and the thin-provisioned volume was automatically set offline because there was not enough free pool space for the volume reserve to grow. Note: To correct this condition, add more space to the pool (such as by adding another member to the pool).
offline-member-down	Volume was automatically set offline because a member that contains volume data is down or cannot be reached by the other members over the network. Note: To correct this condition, fix the problem that caused the member to go down or become unreachable.
offline-lost-blocks	Volume was automatically set offline because blocks were lost. Note: To correct this condition, click the status link and choose how to handle the lost blocks. See <i>Handling Lost Blocks</i> on page 5-49 for more information.

To display details about a specific volume, double-click the volume name in the Volumes panel. You can also click:

Volumes > *volume_name* > Status tab

The Volume Status window appears (Figure 5-7). The Volume and Snapshot Space panel looks slightly different for regular volumes than for thin-provisioned volumes; a thin-provisioned volume example is shown.

Figure 5-7: Volume Status



The General Volume Information panel displays additional information about the volume, including whether iSCSI access to the volume is restricted (only authorized hosts can access the volume) or unrestricted (all hosts can access the volume), or whether there are no access control records for the volume and, therefore, no host access to the volume.

The Volume and Snapshot Space panel shows the reported size, volume reserve, thin provisioning status (enabled or not), snapshot reserve, number of snapshots, status, replication partner, number of iSCSI connections, how much snapshot reserve space the volume is using in the pool, how volume data is distributed among the pool members, and the volume RAID preference, if any.

Clicking the other window tabs displays additional information about the volume and allows you to manage the volume configuration:

- **Access** – Manage the access control records for the volume and its snapshots. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.
- **Snapshots** – Manage volume snapshots. See Chapter 6, *Snapshot Management* for more information.
- **Replication** – Manage volume replication. See Chapter 7, *Basic Replication* for more information.
- **Collections** – Manage volume collections, which enable you to create snapshots or replicas of multiple volumes at the same time. See *Managing Volume Collections* on page 5-15 for more information.
- **Schedules** – Manage snapshot and replica schedules. See *Managing Schedules for Snapshots and Replicas* on page 5-41 for more information.
- **Connections** – Manage the iSCSI connections to the volume and the amount of data read and written. You can select the volume's iSCSI target name in the `iSCSI target` field and copy it to the clipboard, if desired.

Managing Volume Collections

A volume collection consists of one or more volumes from any pool and simplifies the creation of snapshots and replicas. Volume collections are useful when you have multiple, related volumes. In a single operation, you can create snapshots of the volumes (a snapshot collection) or replicas of the volumes (a replica collection).

Volume collection management topics include:

- *Creating a Volume Collection* on page 5-15
- *Displaying Volume Collection Status* on page 5-17
- *Modifying a Volume Collection* on page 5-22
- *Deleting a Volume Collection* on page 5-23


Creating a Volume Collection

To create a volume collection, click:

Volume Collections > Create volume collection

The Create Volume Collection - General window appears (Figure 5-8).


Figure 5-8: Create Volume Collection - General



Specify a name for the collection and an optional description. Then, click **Next**.

The Create Volume Collection - Components window appears (Figure 5-9).


Figure 5-9: Create Volume Collection - Components



Select up to eight volumes for a collection. A volume collection can contain volumes from different pools. Click **Next**.

The Create Volume Collection - Summary window appears (Figure 5-10).

Figure 5-10: Create Volume Collection - Summary



If the configuration is correct, click **Finish** to create the volume collection. To make changes, click **Back**.


Once you create a volume collection, you can create snapshots or replicas:

- For information about creating schedules for a volume collection, see *Managing Schedules for Snapshots and Replicas* on page 5-41.
- For information about creating snapshots of the volumes in a volume collection, see *Creating a Snapshot Collection* on page 6-17.
- For information about creating replicas of the volumes in a volume collection, see *Creating Replicas* on page 7-24.

Displaying Volume Collection Status

To display the volume collections in a group, click **Volume Collections** in the far left panel. The Volume Collections window appears (Figure 5-11).

Figure 5-11: Volume Collections



The Volume Collections panel shows the volumes that comprise each volume collection, including the volume pool, size, snapshot reserve, number of snapshots, status, and replication partner.

Select a collection name or volume name to display related tasks in the Activities panel.

To display details about a specific volume collection, double-click a collection in the Volume Collections panel or click:

Volume Collections > *collection_name* > Status tab

The Volume Collection - Status window appears (Figure 5-12), displaying the snapshot and replication status of the collection and details about the collection volumes.

Figure 5-12: Volume Collection - Status

The screenshot shows the Dell EqualLogic Group Manager interface. The main window is titled "Volume collection Marketing". The left sidebar shows a tree view with "Marketing" selected. The main content area has tabs for "Status", "Snapshots", "Replicas", and "Schedules". The "Status" tab is active, showing "Collection Status" and "Collection Volumes".

Collection Status

Snapshots
 Snapshot collections: 2
 Last snapshot collection: 4/11/07 3:53:13 PM

Replication
 Collection is not replicated

Description: Marketing volumes

Collection Volumes

Total volumes in the collection: 3


Volume name	Storage pool	Reported size	Volume reserve	Snapshot reserve	Volume status	Replicator partner	Number of snapshots	iSCSI connections
exchange...	default	100.0 ...	29.82 ...	29.82 GB	online	group2	12	0
marketing	default	20.01 ...	20.01 ...	20.01 GB	online	group2	4	0
sqlvol	default	23.01 ...	23.01 ...	23.01 GB	online	group2	4	0

The bottom status bar shows "1 outstanding alarm (1 new)".

In the Collection Volumes panel, select a volume name to display related tasks in the Activities panel.

Click the Snapshots tab to display the Volume Collection - Snapshots window (Figure 5-13), which shows the snapshot collections for the volume collection.


Figure 5-13: Volume Collection - Snapshots



See *Managing Snapshot Collections* on page 6-16 for more information about snapshot collections.

Click the **Replicas** tab to display the Volume Collection - Replicas window (Figure 5-14), which shows the replica collections for a volume collection.


Figure 5-14: Volume Collection - Replicas



See Chapter 7 for more information about replica collections.

Click the **Schedules** tab to display the Volume Collection - Snapshot and Replication Schedules window (Figure 5-15).

Figure 5-15: Volume Collection - Schedules



For information about setting up a schedule for a volume collection, see *Managing Schedules for Snapshots and Replicas* on page 5-41.

Modifying a Volume Collection

You can modify a volume collection name, description, and the volumes that comprise the collection.

To modify a volume collection, click:

```
Volume Collection > collection_name > Modify volume collection
```

A dialog box similar to Figure 5-8 is displayed, enabling you to modify the volume collection name and description.

Click the **Components** tab to display a dialog box similar to Figure 5-9, enabling you to select the volumes for the collection.

Click **OK** when you are done.

Deleting a Volume Collection

To delete a volume collection, click:

```
Volume Collection > collection_name > Delete volume collection
```

Confirm that you want to delete the collection.

Note: Deleting a volume collection does not delete the volumes in the collection or any snapshots or replicas associated with the volumes.

Managing Access Controls for Volumes and Snapshots

In a PS Series group, access control records are used to control which hosts can access a volume and its snapshots. It is recommended that you restrict access to prevent multiple hosts from simultaneously accessing a volume or snapshot and possibly corrupting data.

A volume and its snapshots share a single 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 may want to give one host access to both the volume and its snapshots and give another host access only to the volume's snapshots.

In each record, you can specify an IP address, iSCSI initiator name or CHAP user name (or any combination of the three). To access a volume or snapshot, a host must *exactly* match at least one access control record.

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

You can also create a record that allows unrestricted host access. However, this is *not* recommended unless you are testing host access to the volume.

Note: If you want to be certain that no host can access a volume or snapshot, delete *all* its access control records. You can also set the volume or snapshot offline.

Access control record management topics include the following:

- *Creating Access Control Records* on page 5-24
- *Displaying Access Control Records* on page 5-26
- *Modifying an Access Control Record* on page 5-26
- *Deleting an Access Control Record* on page 5-27


Creating Access Control Records

To create an access control record, click:

Volumes > *volume_name* > Access tab

The Volume Access window appears (Figure 5-16), displaying the access control records for the volume and its snapshots.

Figure 5-16: Volume Access



In the Access Control List panel, click **Add**. The Add Access Control Record dialog box appears (Figure 5-17).

Figure 5-17: Add Access Control Record

Create access control record

☐ Authenticate using CHAP user name:

☒ Limit access by IP address (asterisks allowed):
111.222.33.4

☐ Limit access to iSCSI initiator name:

Apply to

☒ volume ☒ snapshots

OK Cancel

Java Application Window

For unrestricted access (not recommended), do not select any access methods.

To restrict access, select the access method and specify the required information:

- **Authenticate using CHAP user name** – Restricts access to hosts that supply the specified CHAP user name (and its associated password or “secret”). The user name must match a local CHAP account or an account on an external RADIUS server, as described in *Using CHAP to Control iSCSI Access to Volumes* on page 3-25.
- **Limit access by IP address** – Restricts access to hosts with the specified initiator IP address (for example, 12.16.22.123). Use asterisks for “wildcards,” if desired (for example, 12.16.*.*). An asterisk can replace an entire octet, but not a digit within an octet.
- **Limit access to iSCSI initiator name** – Restricts access to hosts with the specified iSCSI initiator.

Note: If using IP addresses or iSCSI initiator names to restrict access, be sure to create an access control record for each IP address or iSCSI initiator name presented by an authorized host.

In addition, select whether you want the access control record to apply to the volume, its snapshots, or both (the default).

After entering the access control information, click **OK** to create the record.

Displaying Access Control Records

To display the access control records for a volume and its snapshots, click:

```
Volumes > volume_name > Access tab
```

The Volume Access window appears (Figure 5-16), displaying the records that have been created for the volume and its snapshots.

A host must exactly match at least one of the access control records in order to access the volume or snapshot.

Select a record in the Access Control List panel to display details.

Note: An asterisk in a column means that any value is accepted. If any access control record shows asterisks in the User name, IP address filter, and iSCSI initiator filter columns, the volume or snapshot is unrestricted and *all* hosts can access it.

Modifying an Access Control Record

To modify an access control record, click:

```
Volumes > volume_name > Access tab
```

The Volume Access window appears (Figure 5-16).

In the Access Control List panel, select the record and click **Modify**.

The Modify Access Control Record dialog box appears with the same fields as the Add Access Control Record dialog box (Figure 5-17). Modify the access control record information and click **OK**.

Notes: If you modify the CHAP user name for an access control record or delete an access control record that includes a CHAP user name, hosts that previously matched the record and have a current iSCSI connection to the target (volume or snapshot) may not be able to log in to the target again. It is recommended you modify a CHAP user name in an access control record or delete a record that includes a CHAP user name only when the target has no iSCSI connections.

If you create an access control that limits host access to an iSCSI initiator name, you cannot change the initiator name. Instead, you must delete the record and recreate it using the desired name.

Deleting an Access Control Record

To delete an access control record, click:

Volumes > *volume_name* > Access tab

The Volume Access window appears (Figure 5-16).

In the Access Control List panel, select the record and click **Delete**. Confirm that you want to delete the record.

Note: If you delete all the access control records for a volume or snapshot, there will be no host access to the volume or snapshot.

Connecting Hosts to Volumes and Snapshots

A volume is seen on the network as an iSCSI target, which can only be accessed by an iSCSI initiator installed on a host. Both hardware and software initiators are available from a variety of vendors. Install and configure initiators using the instructions provided by the initiator vendor.

Notes: It is recommended that you visit the Technical Support website and check if there is information related to your iSCSI initiator.

To access a volume, a host must exactly match at least one of the volume's access control records, unless the volume has unrestricted access. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.

Host access to volumes is through TCP/IP port 3260 (the standard iSCSI port). Some iSCSI initiators require this port number.

The exact procedure for logging in to an iSCSI target depends on the iSCSI initiator. See your initiator documentation for details. In most cases, specify the group IP address as either the target portal or the discovery address in the initiator configuration utility.


If the initiator supports the discovery process, it will return a list of iSCSI targets that the host can access. If you are using iSNS, the initiator will automatically discover targets from the iSNS server if the volumes and snapshots are online.

Otherwise, you must specify the volume's target name in the initiator configuration utility. To obtain the iSCSI target name for a volume, click:

Volumes > *volume_name* > Connections tab

The Volume Connections window appears (Figure 5-18) with the iSCSI target name in the iSCSI target field. You can copy it to the clipboard, if desired.

Figure 5-18: Volume Connections



Use the initiator configuration utility to connect (log in) to the iSCSI target for the desired volume. Enter CHAP credentials at this time if required for volume access.

Once the initiator logs in to the iSCSI target, the volume is seen by the host as a normal disk that can be formatted using the usual operating system utilities. For example, you can partition the disk and create a file system.

Note: Some file systems require that a volume or snapshot have read-write permission even if the file system will be read-only. See *Setting Volume Permission* on page 5-48 and *Setting Snapshot Permission* on page 6-12 for information about setting read-write permission on a volume or snapshot.

The Volume Connections window (Figure 5-18) and the Snapshot Status window (Figure 6-4) display the iSCSI connections to volumes and snapshots, including information about the initiator address, connection time, and the amount of data transferred.

Modifying Volume Settings

As needed, you can modify volume settings. See the following topics:

- *Moving a Volume Across Pools* on page 5-29
- *Increasing the Size of a Volume* on page 5-31
- *Decreasing the Size of a Volume* on page 5-33
- *Modifying a Volume Name, Description, and Alias* on page 5-33
- *Enabling and Disabling Thin Provisioning* on page 5-34
- *Setting a Volume RAID Preference* on page 5-35
- *Binding a Volume to a Specific Member* on page 5-37

Moving a Volume Across Pools

By default—or if only one pool exists—a volume will be assigned to the default pool. In multi-member groups, you can separate group space into multiple storage pools and assign members and volumes to the pools. All volume data will be restricted to the members that make up the pool to which the volume is assigned.

As needed, you can move volumes across pools, with no impact on users or applications.

See Chapter 4, *Storage Pool Management* for information about setting up pools.

Notes: Only group administrators can move volumes across pools.

You can move a volume to a different storage pool only if the new pool has free space equal to the size of the volume and any snapshot reserve and local replication reserve for the volume. If the pool does not have enough space, the volume will not move.

While a volume move is in progress, space equal to the volume reserve is allocated in both groups.

If you delete a pool that contains volumes, the volumes will be automatically assigned to the default pool.


If you cancel an in-progress volume move, the volume data will be moved back to the original pool.

To move a volume from one pool to another, click:

Volumes > *volume_name* > Modify volume settings > General tab

The Modify Volume Settings - General dialog box appears (Figure 5-19).

Figure 5-19: Modify Volume Settings - General



Under Storage pool assignment, select the desired pool for the volume.

The Storage pool space table changes to show the new amounts of free and used space in each pool. If the addition of the volume will exceed the capacity of the pool, the table cell showing free pool space will turn red.

Click OK when you are done.

Note: You can cancel an in-progress volume move operation, if desired, and return volume data to the original storage pool. To cancel a volume move operation, click:

Volumes > *volume_name* > Cancel volume moving

Increasing the Size of a Volume

You can increase the size (reported size) of a volume when more space is needed for data.

Note: To recognize the new volume size, hosts that are currently connected to a volume *must* support file system expansion. Not all operating systems, file systems, and applications support increasing the size of a volume or behave in a predictable manner when a volume size is increased. Before increasing a volume size, be sure to understand how the change affects the operating system, file system, or application using the volume.

Be sure you have sufficient free pool space before increasing the volume size.


Because reserved snapshot and replication space is based on a percentage of the volume reserve, increasing the size of a volume usually requires more snapshot and replication space. (For a standard volume, the volume reserve is the same as the reported size; for a thin-provisioned volume, the volume reserve is less than or equal to the reported size.)

For a thin-provisioned volume, increasing its size also increases the in-use warning value and maximum in-use space value (but not the percentages) because they are based on the volume size.

To increase the size of a volume, click:

Volumes > *volume_name* > Modify volume settings > Space tab

The Modify Volume Settings - Space dialog box appears (Figure 5-20).

Figure 5-20: Modify Volume Settings - Space

In the Volume size field, specify the new size. For a thin-provisioned volume, you can also modify the in-use warning value and maximum in-use space value using the slider bars.

The values in the pool space table are updated when you specify a new volume size. See Table 5-1 on page 6 for a description of the values and the effect of changing the volume size.

If the volume is configured for replication, a table showing the replication partner delegated space appears below the pool space table.

If the new volume size exceeds the capacity of the pool or the delegated space on the partner, the table cell showing the new space value space will turn red.

Specify the new volume size and click **OK**.

Decreasing the Size of a Volume

You can decrease the reported size of a volume. This is called shrinking a volume.

Notes: Be careful when decreasing the size of a volume. It is possible to decrease volume space to less than what is currently being used.

Snapshot reserve, local replication reserve, and replica reserve are based on a percentage of the volume reserve. (For a standard volume, the volume reserve is the same as the reported size; for a thin-provisioned volume, the volume reserve is less than or equal to the reported size.) Therefore, when a volume size is decreased, these reserves may also be decreased, unless the reduction will result in the deletion of snapshots. In this case, the snapshot reserve and local replication reserve will be set to a percentage that is sufficient to retain the existing snapshots. An event message will be generated if this occurs.

To decrease the size of a volume, you must use the Group Manager command line interface (CLI). See the `volume shrink` command in the *CLI Reference* for more information.

Modifying a Volume Name, Description, and Alias

You can modify the name that is used to identify the volume for administrative purposes and the volume description.

You can also modify the iSCSI target name alias (public name) for a volume. If you are using the default group-wide volume setting, the alias is the same as the volume name. The volume alias is seen by some iSCSI initiators along with the iSCSI target name. It can help administrators identify the volume.

Note: Of you rename a volume, the iSCSI target name (and any existing snapshot or replica set names) will not change. However, if you rename a volume and the volume alias is set to be the same as the volume name, the alias will change accordingly.

To modify a volume name or description, click:

Volumes > `volume_name` > Modify volume settings > General tab

The Modify Volume Settings - General dialog box appears (Figure 5-19).

Modify the name in the `Volume name` field. The volume name must be unique and can contain up to 64 alphanumeric characters (including periods, hyphens, and colons).

Optionally, specify a description for the volume in the `Description` field.

When you are finished making changes, click `OK`.

To modify the alias for a volume, click:

Volumes > `volume_name` > Modify volume settings > Advanced tab

The Modify Volume Settings - Advanced window appears (Figure 5-21).

Specify a volume alias in the `Public alias` field and click `OK`.

Enabling and Disabling Thin Provisioning

You can enable or disable thin provisioning on an existing volume.

Disabling thin provisioning on a volume will fully allocate the volume size from free pool space (the reported size and the volume reserve will be the same). Enabling thin provisioning on a volume decreases the amount of space allocated to the volume (the reported size will be greater than the volume reserve).

Because snapshot reserve, local replication reserve, and replica reserve space are all based on a percentage of the volume reserve, enabling and disabling thin provisioning will also affect these values.

Note: Be careful when enabling or disabling thin provisioning on a volume. Enable this feature only for volumes that you have monitored and determined to be underused; that is, the amount of in-use space stays significantly less than the reported size. Consider disabling thin provisioning for volumes that are nearing their maximum in-use space value.

See *Understanding Thin Provisioning* on page 5-8 for details on thin provisioning.

When *enabling* thin provisioning on a volume, be aware of these issues:

- Unused space that was allocated to the volume is returned to free pool space.
- The volume reserve will decrease. Because the snapshot reserve, local replication reserve, and replica reserve are based on a percentage of the volume reserve, they will also decrease, unless the reduction will result in the deletion of snapshots. In this case, the snapshot reserve and local replication reserve will be set to a percentage that is sufficient to retain the existing snapshots. An event message will be generated if this occurs.

When *disabling* thin provisioning on a volume, be aware of these issues:

- The volume will require additional volume reserve, snapshot reserve, local replication reserve, and replica reserve, unless the thin-provisioned volume was fully allocated (that is, the volume reserve and reported size were the same).
- If the volume is configured for replication, ensure that the partner has sufficient delegated space to handle the increase in replica reserve.
- In some cases, disabling thin provisioning on a volume will automatically decrease the snapshot reserve *percentage* if disabling thin provisioning would result in a large amount of unused space reserved for snapshots.

This can occur if the volume was originally converted to a thin-provisioned volume, and the conversion resulted in an increase in the snapshot reserve percentage to prevent snapshots from being deleted. (For example, to maintain the existing snapshots, the snapshot reserve percentage could have been increased from 100% to 500%.) If thin provisioning is later disabled on the volume, and the snapshot reserve percentage would result in an excessive amount of snapshot reserve, the percentage will be reduced to a lower, more appropriate value, closer to 100%, that can still retain the existing snapshots.

To enable or disable thin provisioning on a volume, click:

Volumes > *volume_name* > Modify volume settings > Space tab

The Modify Volume Settings - Space dialog box appears (Figure 5-20).

Select or deselect the checkbox next to `Thin provisioned volume` to enable or disable thin provisioning on the volume.

Optionally, if enabling thin provisioning, make any desired adjustments to the minimum volume reserve, in-use space warning value, or maximum in-use space value by moving the relevant sliders. You cannot set the maximum value below the warning value.

When you are finished, click `OK`.

Setting a Volume RAID Preference

A PS Series group uses performance load balancing metrics (enabled by default) to automatically identify the RAID level that will provide optimal performance for the volume and attempts to locate the data on pool members configured with that RAID level.

For example, if the group determines that a volume performs best with RAID 10, it will attempt to place the volume data on pool members configured with RAID 10.

However, in some cases, a group administrator may want to manually specify a RAID preference for a volume (either RAID 10, RAID 50, or RAID 5). If a volume is assigned a RAID preference, the group will *attempt* to place volume data on pool members with that RAID level. Setting a RAID preference will override automatic performance load balancing for the volume.


Note: Volume data is restricted to the members in the storage pool to which the volume is assigned. If you attempt to set a RAID preference for a volume and the RAID level does not exist on any of the pool members, or if you delete the only pool member with the preferred RAID level or change its RAID level, you will be notified that the volume's RAID preference will not be honored (enforced) until a member is configured with the preferred RAID level. If this occurs, only automatic capacity load balancing across the pool members will be enforced for the volume.

See *Displaying Storage Pool Configuration and Status* on page 4-6 for information about identifying the RAID levels of the pool members.

To set a RAID preference for a volume, click:

Volumes > *volume_name* > Modify volume settings > Advanced tab

The Modify Volume Settings - Advanced dialog box appears (Figure 5-21).

Figure 5-21: Modify Volume Settings - Advanced

Under Volume storage preference, select the desired RAID level and click **OK**.

Binding a Volume to a Specific Member

You can configure a volume so that the group stores the volume data on one specific pool member, instead of allowing the volume data to be distributed across multiple members. If you restrict a volume to one member, you will not obtain performance benefits from load balancing data and network connections across multiple members.

Note: Binding a volume to a specific member will override automatic performance and capacity load balancing and any RAID preference for the volume. See *Controlling Performance Load Balancing* on page 3-38 and *Setting a Volume RAID Preference* on page 5-35 for more information.

You can bind a volume only to a member that is in the pool to which the volume is assigned. In addition, if you bind a volume to a member and then delete that member from the pool or group, the volume bind will be cancelled.

To bind a volume to a member, you must use the Group Manager command line interface (CLI). You cannot use the GUI. See the PS Series *CLI Reference* manual for more information.

Cloning Volumes

Cloning a volume creates a new volume with a new name and iSCSI target, but the same size, contents, and thin provisioning setting as the original volume at the time of the cloning. The new volume is located in the same pool as the original volume and is available immediately. Cloning a volume does not affect the original volume, which continues to exist after the cloning operation.

A cloned volume consumes 100% of the original volume size from free space in the pool in which the original volume resides. If you want to create snapshots or replicas of the new volume, you will need additional pool space.

To clone a volume, click:

```
Volumes > volume_name > Clone volume
```

The Clone Volume - General Settings dialog box (Figure 5-22) appears.

Figure 5-22: Clone Volume - General Settings

Clone volume vol1

Step 1 - Volume Settings

General properties:

* Volume name: clonevol1

Description: cloned from vol1

Storage pool assignment:

Storage Pool	Capacity	Free	Drives
☒ default	5.23 TB	3.04 TB	SATA

Back Next Finish Cancel

Java Application Window

Specify information about the new volume in the following fields:

- **Volume name** – Unique name, up to 64 alphanumeric characters (including periods, hyphens, and colons), used to identify the volume for administrative purposes. Host access to the volume is always through the iSCSI target name, not the volume name.
- **Description** – Optional volume description.

Note: You cannot choose a different storage pool (if available) during the clone operation.

Click **Next**.

The Clone Volume - Space Reserve dialog box appears (Figure 5-23).

Figure 5-23: Clone Volume - Space Reserve

Clone volume vol1

Step 2 - Space reserve

Volume space:

* Volume size: ☐ MB ☒ GB (max: 15.0 TB)

☒ Thin provisioned volume (use sliders below to adjust settings)

Snapshot space:

* Snapshot reserve: % of volume reserve

Reported volume size 10.0 GB

Free: 10% | Unreserved 8.99 GB: 60% | 80%

Estimated changes in storage pool default

	Current	New	Change
Volume reserve	290.41 GB	291.42 GB	1.01 GB
Snapshot reserve	290.41 GB	291.42 GB	1.01 GB
Replication reserve	290.41 GB	290.41 GB	0 MB
Delegated space	1.32 TB	1.32 TB	0 MB
Free pool space	3.04 TB	3.03 TB	-2.02 GB

Back Next Finish Cancel

Java Application Window

Note: You cannot change the volume size during the clone operation. You can change the size and other volume attributes after creating the volume, as described in *Modifying Volume Settings* on page 5-29.


In the **Snapshot reserve** field, optionally, specify the amount of space, as a percentage of the volume reserve, to reserve for snapshots of the new volume. The default is the group-wide volume setting. Snapshot space is reserved even if you never create a snapshot. You cannot create snapshots of a volume until you reserve space.

The values in the pool space table shown in Figure 5-23 reflect the effect of the new volume on the pool space. See Table 5-1 on page 6 for a description of the values.

If cloning the volume will exceed the capacity of the pool or the delegated space on the partner, the relevant table cell will turn red.

Click **Next** to display the Clone Volume – iSCSI Access Policy dialog box (Figure 5-24).


Figure 5-24: Clone Volume - iSCSI Access Policy



This dialog box enables you to create an access control record that restricts host access to the new volume and its snapshots, as described in *Creating Access Control Records* on page 5-24.

Select the method of restricting access and specify the required information.

Click **Next** to display the Clone Volume – Summary dialog box, which summarizes the volume configuration (Figure 5-25).

Figure 5-25: Clone Volume - Summary

If the configuration is correct, click **Finish** to clone the volume. Click **Back** to make changes.

Managing Schedules for Snapshots and Replicas

Schedules enable you to create a snapshot or replica of a volume or a volume collection at a specific time or on a regular basis (for example, hourly or daily).

Using a schedule can result in a large number of snapshots or replicas, so be sure that you have sufficient space. In addition, you can place a limit on the maximum number of snapshots or replicas created from a schedule. See *Modifying Snapshot Settings* on page 6-3 for information on snapshot space. See *Understanding Replication Space Requirements* on page 7-5 for information on replication space.

Notes: Be careful not to configure a snapshot or replication schedule on a volume that is also part of a collection with a schedule, as the schedules might interfere with each other.

If a volume is configured for manual transfer replication, do not create a schedule for the volume until the first replication completes.

To create a snapshot or replica schedule for an individual volume, click:

Volumes > *volume_name* > Create schedule

To create a snapshot or replica schedule for a volume collection, click:


Volume Collections > *collection_name* > Create schedule

The Create Schedule - Schedule Type dialog box appears (Figure 5-26).

Specify the following information:

- Schedule name.
- Whether the schedule will be used for snapshots or replication.
- Type of schedule, either once, hourly, or daily. You can also choose to use an existing schedule and modify it to meet your needs by selecting *Reuse existing schedule*.
- Whether you want the schedule to be enabled (actively running) or disabled.

Figure 5-26: Create Schedule - Schedule Type




Click **Next** to display a dialog box that allows you to enter the parameters for the chosen type of schedule.

Depending on the type of schedule, parameters can include the following:


- Dates that the schedule will start and, optionally, stop running.
- How often to run the schedule.
- Time to create the snapshots or replicas.
- Frequency of snapshots or replicas.
- Maximum number of snapshots or replicas created by the schedule that the group will retain. If you reach the maximum, the oldest snapshots or replicas created by the schedule will be automatically deleted before a new snapshot or replica is created. If a snapshot to be deleted has active iSCSI connections, the connections will be terminated before the snapshot is deleted.
- Whether the snapshots will have read-write permission.

For example, if you select **Run once**, the Create Schedule - Run Once dialog box appears (Figure 5-27). Note that the schedule dialog boxes shown below refer to setting up a snapshot schedule. Similar dialog boxes appear when setting up a replica schedule.


Figure 5-27: Create Schedule - Run Once



If you select **Hourly** schedule, the Create Schedule - Hourly Schedule dialog box appears (Figure 5-28).


Figure 5-28: Create Schedule - Hourly Schedule

If you select `Daily` schedule, the Create Schedule - Daily Schedule dialog box appears (Figure 5-29).

Figure 5-29: Create Schedule - Daily Schedule

After specifying the schedule parameters, click **Next**. The Create Snapshot Schedule - Summary dialog box appears (Figure 5-30).

Figure 5-30: Create Schedule - Summary




If the schedule configuration is correct, click **Finish** to create the schedule. Click **Back** to make changes.

To display schedules for a specific volume, click:

Volumes > *volume_name* > Schedules tab

The Volume Schedules window appears (Figure 5-31).

Figure 5-31: Volume Schedules



The Schedule Summary panel shows the status of the volume schedules and the creation time of the next scheduled snapshot or replica. A running schedule is a schedule that is set up to create snapshots or replicas.


In the Snapshot and Replication Schedules panel, select a schedule to display schedule details and relevant tasks in the Activities panel. Select **Also display schedules for collections that include the volume** to also display schedules for volume collections that include the volume.

To display the schedules for a volume collection, click:

Volume Collections > *collection_name* > Schedules tab

The Volume Collection Schedules window appears (Figure 5-32).

Figure 5-32: Volume Collection Schedules



Select a schedule to display schedule details and schedule-related tasks in the Activities panel.

To modify a schedule, select the schedule and click **Modify**. Dialog boxes appear (the same as when creating a schedule), enabling you to modify the schedule parameters.

Note: You cannot change the name of a schedule, whether the schedule is for snapshots or replicas, or the type of schedule (one time, daily, or weekly).

To disable a schedule, select the schedule and, in the Activities panel, click **Disable schedule**. To enable a schedule, select the schedule and click **Enable schedule**.

To delete a schedule, select the schedule and click **Delete**. Confirm that you want to delete the schedule.

Note: Deleting a schedule will not delete any snapshots or replicas created from the schedule.

Setting a Volume Offline or Online

You can set a volume offline to make it temporarily inaccessible to hosts (iSCSI initiators). Current iSCSI connections will be lost. Offline volumes cannot be found by iSCSI initiators during discovery or through iSNS.

To make the volume accessible again, set the volume online.

To set a volume offline, click:

```
Volumes > volume_name > Set volume offline
```

To set a volume online, click:

```
Volumes > volume_name > Set volume online
```

Once you set a volume online, hosts can connect to it if they have the right access credentials, as described in *Connecting Hosts to Volumes and Snapshots* on page 5-27.

Note: If all the group members that contain volume data are not online, you cannot set the volume online.

Setting Volume Permission

A volume can have either read-write (default) or read-only permission.

Note: To change the permission of an online volume to read-only, you must first set the volume offline.

To set a volume to read-only, click:

```
Volumes > volume_name > Set read-only
```

Note: An application may fail if it attempts to write to a volume that has read-only permission.

To set a volume to read-write, click:

```
Volumes > volume_name > Set read-write
```

Deleting Volumes

You can delete a volume and return the disk space to the free group space.

Notes: If you delete a volume, all the volume snapshots are also deleted. However, volume replicas on the secondary group are not affected.

Active iSCSI connections to a volume will be terminated before the volume is deleted.

To delete a volume, click:

```
Volumes > volume_name > Delete volume
```

Confirm that you want to delete the volume.


Handling Lost Blocks

In rare circumstances, a volume or snapshot may experience lost blocks. For example, if a power failure is followed by the failure of a control module cache battery (resulting from more than 72 hours of battery operation after the power failure or other factors), lost blocks may occur.

If volume blocks have been lost, the Volume Summary window (Figure 5-6) will display the current status of the volume as `offline-lost-blocks`. If snapshot blocks have been lost, the Volume Snapshots window (Figure 6-3) will display the current status of the snapshot as `offline-lost-blocks`.

In addition, an event message will be generated indicating the lost blocks.

To handle lost blocks, click the `offline-lost-blocks` link. The Acknowledge lost blocks dialog box appears (Figure 5-33).

Figure 5-33: Acknowledge Lost Blocks

If lost blocks occur, you can do one of the following:

- Set the volume or snapshot online but retain the lost blocks.

If an application attempts to read a lost block, an error will occur. If the block is rewritten, no error will occur, and the block will no longer be marked lost. The member or members containing volume or snapshot data will continue to show a status of `RAID lost blocks` until the blocks have been written with new data.

To retain the lost blocks, select `Set the volume online but retain the lost blocks`. The current status will be changed to `online-lost-blocks`.

- Mark the lost blocks valid and set the volume or snapshot online.

This option may present a security risk, and the blocks may contain old data. The status will be cleared for the member or members containing volume or snapshot data.

To mark the lost blocks valid, select `Mark the lost blocks valid and set the volume online`. The current status will be changed to `online`.

6 Snapshot Management

In a PS Series group, you can create snapshots of volumes. A snapshot enables you to capture the contents of a volume at a specific point in time. Snapshots greatly simplify and improve the performance of backup and recovery operations.

Only group administrators or pool administrators can create or modify volume snapshots. For more information, see *Managing Group Administration Accounts* on page 3-9.

Snapshot topics include the following:

- *Introduction to Snapshots* on page 6-1
- *Modifying Snapshot Settings* on page 6-3
- *Creating Snapshots* on page 6-5
- *Displaying Snapshot Configuration and Status* on page 6-6
- *Setting a Snapshot Online or Offline* on page 6-11
- *Setting Snapshot Permission* on page 6-12
- *Modifying a Snapshot Name, Description, and Alias* on page 6-12
- *Scheduling Snapshots* on page 6-14
- *Restoring a Volume from a Snapshot* on page 6-14
- *Cloning a Snapshot* on page 6-16
- *Deleting a Snapshot* on page 6-16
- *Managing Snapshot Collections* on page 6-16

Introduction to Snapshots

A snapshot is a point-in-time copy of volume data that can protect against mistakes, viruses, or database corruption. Snapshot creation does not disrupt access to the volume.

Like volumes, snapshots appear on the network as iSCSI targets and can be set online and accessed by hosts with iSCSI initiators. You can recover volume data by restoring a volume from a snapshot or by cloning a snapshot, which creates a new volume.

To create a snapshot, you must first reserve snapshot space for the volume. Snapshot reserve is consumed from the same pool in which the volume resides, and is based on a percentage of the current volume reserve.

There are several ways to create snapshots:

- You can create a snapshot of an individual volume at the current time, or you can set up one or more schedules to automatically create a snapshot at a specific time in the future or on a regular basis.

See *Creating Snapshots* on page 6-5 and *Scheduling Snapshots* on page 6-14 for information.

- You can create a snapshot of each volume in a volume collection at the current time, or you can set up one or more schedules to automatically create a snapshot of each volume in the collection at a specific time in the future or on a regular basis. When you create snapshots of the volumes in a collection, the resulting set of snapshots is called a snapshot collection.

See *Creating a Snapshot Collection* on page 6-17 and *Scheduling Snapshots* on page 6-14 for information.

- You can create snapshots of multiple volumes at the current time. The resulting set of snapshots is called a custom snapshot collection, because it is not associated with a volume collection.

See *Creating a Custom Snapshot Collection* on page 6-18 for information.

A snapshot name is generated automatically, based on the volume name and the date and time when the snapshot was created (for example, `staff1-2006-05-08-11:29:29.1`). However, where the GUI displays snapshots under a volume name (for example, in the far left panel), the snapshots are identified only by the timestamp.

An iSCSI target name for a snapshot is also automatically generated. The iSCSI target name will consist of the iSCSI prefix (`iqn.2001-05.com.equallogic`), followed by a generated string, followed by the snapshot name.

As with volumes, host access to snapshots is controlled by access control records. A volume and its snapshots share a list of access control records. You can specify that a record apply only to the volume, only to the volume's snapshots, or to both the volume and its snapshots. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.

PS Series group snapshot technology uses a copy-on-write technique. When users make changes to the base volume (or to the snapshot), the group tracks the changes in the snapshot reserve.

A warning is issued when the amount of free snapshot reserve for the volume falls below the user-defined warning threshold. You can increase the reserve, as needed, if pool space is available.

Note: A snapshot depends on its base volume. If a volume goes offline because of a problem (for example, because a member is offline due to network problems), all its snapshots will be set offline. If you delete a volume, all its snapshots are also deleted.

There are several mechanisms that control how many volume snapshots can exist at one time:

- Snapshot reserve. The size of the snapshot reserve always restricts the number of snapshots for a volume, regardless of any schedule setting for the maximum number to keep. In some cases, you may never reach the maximum number of snapshots to keep because the snapshot reserve is too small to hold that many snapshots.
- Snapshot space recovery policy. When the reserved snapshot space has been exceeded, the group will delete the oldest snapshots or set the volume and its snapshots offline, depending on the volume's snapshot space recovery policy setting.
- Maximum snapshots to keep schedule setting. When you create a snapshot schedule, you must specify the maximum number of snapshots created by the schedule to keep in the snapshot reserve. Once you reach the maximum, the oldest snapshots created by the schedule will be automatically deleted before a new snapshot is created.

In some cases, you may want to preserve the data in a snapshot. To do this, you can clone the snapshot, as described in *Cloning a Snapshot* on page 6-16. For example, if a snapshot will be automatically deleted because of lack of snapshot reserve, you can clone the snapshot to preserve the data.

Modifying Snapshot Settings

There are three snapshot settings for a volume:

- Snapshot reserve size
- Low-snapshot space warning threshold
- Policy to follow when snapshot space is exceeded


When creating a volume, the default snapshot settings are the group-wide settings, as described in *Creating Volumes* on page 5-3. You can change the snapshot settings for an existing volume, as needed.

To modify the snapshot settings for a volume, click:

Volumes > *volume_name* > Modify snapshot settings

The Modify Snapshot Settings dialog box appears (Figure 6-1).

Figure 6-1: Modify Snapshot Settings



You can change the following settings:

- **Snapshot reserve.** Specify the amount of space, as a percentage of the volume reserve, to reserve for snapshots of the volume.

When you change the snapshot reserve, the values in the pool space table shown in Figure 6-1 are updated:

- Volume reserve. Pool space allocated to the volume.
- Snapshot reserve. Pool space reserved for snapshots of the volume.
- Free pool space. Unused pool space.

If increasing the snapshot reserve will exceed the capacity of the pool, the new free pool space cell will turn red.

Note that if you decrease the snapshot reserve below the amount of space that is currently holding snapshot data, the snapshot space recovery policy will be enforced, as described below.

- **Snapshot space recovery policy.** Select the action to take when the reserved snapshot space has been exceeded, either delete oldest snapshot or put the volume and its snapshots offline. If a snapshot has active iSCSI connections, the connections will be terminated before the snapshot is deleted or the volume and its snapshots are set offline.
- **Snapshot space warning level.** Change the warning level by moving the slider in the Modify Snapshot Settings dialog box (Figure 6-1). When the in-use snapshot reserve reaches the value specified by the warning level (set as a percentage of the snapshot reserve), an alarm is generated. For example, if a volume is 200 MB with 100% of the volume size reserved for snapshots, and the warning level is 90%, a warning will be generated when in-use snapshot reserve reaches 180 MB.

When you have completed the changes to the snapshot settings, click **OK**.

Creating Snapshots


You can create a snapshot of a volume at the current time. Snapshot creation occurs instantly. The snapshot will be located in the same pool as the base volume.

Before you can create a snapshot of a volume, snapshot space must be reserved for the volume. See *Modifying Snapshot Settings* on page 6-3 for information about reserving snapshot space.

To create a snapshot at the current time, click:

Volumes > *volume_name* > Create snapshot

The Create Snapshot dialog box appears (Figure 6-2).

Figure 6-2: Create Snapshot

You can specify the following information in the Create Snapshot dialog box:

- Optional snapshot description.
- Whether you want the snapshot set offline (default) or online.
- Whether you want the snapshot to be read-write (default) or read-only.

Click **OK** to create the snapshot.

The snapshot will appear in the far left panel, under the volume name, identified by the date and time when the snapshot was created, as shown in Figure 6-3.


Displaying Snapshot Configuration and Status

To display information about a volume's snapshots, click:

Volumes > *volume_name* > Snapshots tab

The Volume Snapshots window appears (Figure 6-3).

Figure 6-3: Volume Snapshots



In the far left panel, snapshots appear under the volume name and are identified by the date and time when the snapshot was created.

The Snapshot Summary panel shows details about the volume's snapshot configuration, including the snapshot settings, the amount of free snapshot reserve, and any snapshot schedules for the volume.

Note: If snapshot reserve is low, consider increasing the size of the reserve, as described in *Modifying Snapshot Settings* on page 6-3. If the snapshot reserve is exceeded, the snapshot space recovery policy will be enforced.

The Snapshots panel shows details about each snapshot, including the date and time it was created, whether the snapshot is part of a snapshot collection, whether the snapshot was created from a schedule, the current status, permission, and iSCSI connections.

Select a snapshot in the Snapshots panel to display related tasks in the Activities panel. For example, you can modify snapshot properties (name, description, or alias), clone the snapshot, set the snapshot online or offline, set the snapshot permission, or delete the snapshot.

Place the mouse over a snapshot to display a pop-up message displaying both the current snapshot status and the requested snapshot status. The requested status is set by an administrator and described in Table 6-1. The Status column in the Volume Snapshots window displays the current snapshot status, as described in Table 6-2.

Table 6-1: Requested Snapshot Status

Status	Description
online	Administrator set the snapshot online. Authorized hosts can access the snapshot.
offline	Administrator set the snapshot offline. Hosts cannot access the snapshot.
online-no-space	Administrator set the snapshot offline, despite lost blocks. Authorized hosts can access the snapshot. If an application attempts to read a lost block, an error will occur. If the block is rewritten, no error will occur, and the block will no longer be marked lost.

Table 6-2: Current Snapshot Status

Status	Description
online	Administrator set the snapshot online. Authorized hosts can access the snapshot, and no failures have occurred. Online snapshots appear in the far left panel in black text.
offline	Administrator set the snapshot offline. Hosts cannot access the snapshot, but no failures have occurred. Offline snapshots appear in the far left panel in gray text.
offline-lost-blocks	Snapshot was automatically set offline because lost blocks have occurred. Hosts cannot access the snapshot. Note: To correct this condition, click the requested status link and choose how to handle the lost blocks. See <i>Handling Lost Blocks</i> on page 5-49 for more information.
offline-inconsistent	Snapshot was automatically set offline due to data inconsistencies. Contact your array support provider for information on how to proceed.

Table 6-2: Current Snapshot Status (Continued)

Status	Description
offline-member-down	Snapshot was automatically set offline because a member that contains snapshot data is down or cannot be accessed by the other members over the network. Hosts cannot access the snapshot. Note: To correct this condition, fix the problem that caused the member to go down or become unreachable on the network.
offline-nospace-auto-grow	Volume reserve was consumed, and the thin provisioned volume and its snapshots were set offline because there was not enough free pool space for the volume reserve to grow. Note: To correct this condition, add more space to the pool (such as by adding another member to the pool).
offline-snap-reserve-met	Snapshot was automatically set offline because the reserved snapshot space has been exceeded and the snapshot recovery policy dictates that the snapshots be set offline. Hosts cannot access the snapshot. Note: To correct this condition, increase the amount of reserved snapshot space.
offline-thin-max-grow-met	Thin-provisioned volume and its snapshots were set offline because in-use space reached the maximum allowed. Note: To correct this condition, increase the value of the maximum in-use space for the volume or increase the volume size, as described in <i>Modifying Volume Settings</i> on page 5-29.


Note: If a snapshot is offline but it was set online by an administrator, ensure that all members that contain snapshot and volume data are online. Also, check the volume status. If a volume is offline but was set online by an administrator, all its snapshots will be set offline.

To display snapshot details, click:

Volumes > *volume_name* > *snapshot_timestamp* > Status tab

The Snapshot Status window appears (Figure 6-4).

Figure 6-4: Snapshot Status




The Snapshot Status window displays details about the snapshot, including whether the snapshot was created from a volume collection or schedule, the current iSCSI connections, and the iSCSI target name. The Activities panel shows snapshot-related tasks.

If the snapshot belongs to a snapshot collection, click [View snapshot collection](#) to display details of the snapshot collection (Figure 6-14 and Figure 6-15).

Click the [Access](#) tab to display the Snapshot Access window (Figure 6-5), which shows the access control records for the snapshot and its base volume. A host must *exactly* match *one* of the access control records in order to access the snapshot.

Figure 6-5: Snapshot Access



Select a record in the Access Control List panel to display details. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for more information.

Setting a Snapshot Online or Offline

An online snapshot can be accessed by hosts. Set a snapshot offline to ensure that no hosts can access the snapshot. If a snapshot was online and you set it offline, any iSCSI connections to the snapshot will be lost.

Notes: If you allow a host to write to a snapshot, you can no longer be certain that the snapshot represents the volume data at the time the snapshot was created. If you set a snapshot online, set its permission to read-only.

To set a snapshot online, click:

Volumes > volume_name > snapshot_timestamp > Set snapshot online

To set a snapshot offline, click:

Volumes > volume_name > snapshot_timestamp > Set snapshot offline

Setting Snapshot Permission

A snapshot can have read-only or read-write permission.

To change the permission of an online snapshot to read-only, you must first set the snapshot offline.

Note: If you set read-write permission on a snapshot and set it online, hosts will be able to access and potentially write to the snapshot. This means the snapshot data is not guaranteed to represent the volume data at the time the snapshot was created.

To set a snapshot to read-only permission, click:

```
Volumes > volume_name > snapshot_timestamp > Set snapshot offline
```

```
Volumes > volume_name > snapshot_timestamp > Set read-only
```

To set a snapshot to read-write permission, click:

```
Volumes > volume_name > snapshot_timestamp > Set read-write
```

Modifying a Snapshot Name, Description, and Alias

You can modify a snapshot name, description, or alias. The alias (public name) is seen by some iSCSI initiators along with the iSCSI target name. It can help administrators identify the snapshot.

Notes: If you rename a snapshot, and the iSCSI target name alias is set to be the same as the snapshot name (as described in *Modifying the Group-Wide Volume Settings* on page 3-36), the alias will change to match the new snapshot name.

The iSCSI target name will *not* change if you rename a snapshot.

To modify a snapshot name, description, or alias, click:

```
Volumes > volume_name > snapshot_timestamp > Modify snapshot properties
```

The Modify Snapshot Properties – General dialog box appears (Figure 6-6).

Figure 6-6: Modify Snapshot Properties - General

Modify the snapshot name and description, as desired.

To modify the iSCSI target name alias for a snapshot, click the `iSCSI` tab. The Modify Snapshot Properties – iSCSI dialog box appears (Figure 6-7).

Figure 6-7: Modify Snapshot Properties - iSCSI

Enter the alias in the `Public alias` field and click `OK`.

Scheduling Snapshots

You can set up a schedule that enables you to create snapshots of a volume or volume collection at a specific time in the future or on a regular basis.

Using a schedule can result in a large number of snapshots, so be sure that you have sufficient snapshot reserve. See *Modifying Snapshot Settings* on page 6-3 for information on reserving snapshot space. You can also place a limit on the maximum number of snapshots created from a schedule.

See *Managing Schedules for Snapshots and Replicas* on page 5-41 for information about creating a schedule.

Restoring a Volume from a Snapshot

You can restore a volume from a snapshot. The restored volume will contain the data that existed in the base volume at the time the snapshot was created and will have the volume's name, iSCSI target name, and thin-provisioning setting. The snapshot will still exist after the restore operation.

After the volume restore completes, you can set the volume online and reconnect iSCSI initiators to the volume. You can also change the volume properties, if desired.

Before the restore operation begins, a snapshot of the current volume state is automatically created. If the restore operation fails, or if for any reason you want to return to the current volume state, you can restore the volume from this snapshot, which is named according to the normal rules for naming snapshots (volume name followed by the date and time the snapshot was created).

Notes: The volume and snapshot must be offline during the restore operation.

All members that contain data from the base volume or the snapshot must be online to restore the volume from the snapshot.

Before restoring a volume from a snapshot, be sure to disconnect any iSCSI initiators from the volume. Otherwise, the restored volume may be corrupted when initiators reconnect to it.

To restore a volume from a snapshot, follow these steps:

1. Disconnect any iSCSI initiators from the volume. Follow the instructions for your operating system and initiator.

2. Set the volume offline by clicking:

Volumes > *volume_name* > Set volume offline

3. Set the snapshot offline by clicking:


Volumes > *volume_name* > *snapshot_timestamp* > Set snapshot offline

4. Restore the volume from the snapshot by clicking:

Volumes > *volume_name* > Restore from snapshot

The Restore volume from snapshot dialog box appears.

Figure 6-8: Restore Volume from Snapshot



5. Select the snapshot from the list and click OK to restore the volume.

Cloning a Snapshot

You can clone a snapshot to create a new volume with the same contents, size, and thin-provisioning setting as the base volume at the time the snapshot was created. The new volume is available immediately in the same storage pool as the base volume and can be used like any other volume. The snapshot will still exist after cloning.

A cloned snapshot consumes 100% of the base volume size from free pool space. If you want to create snapshots or replicas of the new volume, you will need additional space.

To clone a snapshot, click:

```
Volumes > volume_name > snapshot_timestamp > Clone snapshot
```

A series of dialog boxes appear. The information is the same as what is required when you clone a volume. For example, you must specify the new volume's name, snapshot settings, and access controls. See *Cloning Volumes* on page 5-38 for details.

Deleting a Snapshot

To delete a snapshot, click:

```
Volumes > volume_name > snapshot_timestamp > Delete snapshot
```

Confirm that you want to delete the snapshot.

Note: If you delete a snapshot that is part of a snapshot collection or a custom snapshot collection, the collection status will be `incomplete`.

Managing Snapshot Collections

You can create snapshots of multiple volumes at the same time. For example, you can create snapshots of the volumes in a volume collection in one operation. The resulting set of snapshots, one for each volume, is called a snapshot collection and represents the contents of the volumes at the time the snapshots were created.

In addition, you can create snapshots of multiple volumes in a single operation, without using a volume collection. The resulting set of snapshots, one for each volume, is called a custom snapshot collection.

You can create snapshot collections at the current time (as described in *Creating a Snapshot Collection* on page 6-17) or by using a schedule (as described in *Scheduling Snapshots* on page 6-14).

Notes: The snapshots in a snapshot collection are treated the same as any other snapshot and are subject to being automatically deleted, if designated by the volume snapshot recovery policy or the schedule's maximum snapshots setting. If a snapshot that is part of a snapshot collection is deleted, the snapshot collection status will be `incomplete`.

Snapshot collection management topics include the following:

- *Creating a Snapshot Collection* on page 6-17
- *Creating a Custom Snapshot Collection* on page 6-18
- *Displaying Snapshot Collection Status* on page 6-20
- *Modifying a Snapshot Collection* on page 6-23
- *Deleting a Snapshot Collection* on page 6-25

Creating a Snapshot Collection


Before you create a snapshot collection of a volume collection, each volume must have sufficient snapshot space reserved. See *Modifying Snapshot Settings* on page 6-3 for information about reserving snapshot space.

To create a snapshot collection of a volume collection, click:

Volume Collection > `collection_name` > Create snapshot now

The Volume Collection - Create Snapshot dialog box appears (Figure 6-9).

Figure 6-9: Volume Collection - Create Snapshot



Enter an optional description for the snapshot collection and click **OK** to create the snapshot collection.

The snapshot collection will appear under the volume collection name in the far left panel, identified by the date and time when the snapshots were created. The same timestamp is used to identify the individual snapshots, which appear in the GUI under their volume names.

The default snapshot collection name is the volume collection name followed by the date and time when the snapshots were created (for example, `mailcol-2006-06-20-11:17:04.67`). The snapshot collection name appears when you place the mouse over the snapshot collection in the far left panel. The name also appears in the Snapshot Collection window header (Figure 6-14).

Creating a Custom Snapshot Collection

If you do not want to use a volume collection, you can still create snapshots of multiple volumes in a single operation. The resulting set of snapshots, one for each volume, is called a custom snapshot collection.


Before you can create a custom snapshot collection, all the volumes must have sufficient snapshot space reserved. See *Modifying Snapshot Settings* on page 6-3 for information about reserving snapshot space.

Note: If you often create snapshots of the same set of volumes, you may want to create a volume collection, as described in *Managing Volume Collections* on page 5-15, and create a snapshot collection of the volume collection.


To create a custom snapshot collection, click:

Volumes > Multi-volume snapshot


The Create Multi-Volume Snapshot - General Settings dialog box appears (Figure 6-10).

Figure 6-10: Create Multi-Volume Snapshot - General Settings

Enter a name for the custom snapshot collection and an optional description and click **Next**. The Create Multi-Volume Snapshot - Components dialog box appears (Figure 6-11).

Figure 6-11: Create Multi-Volume Snapshot - Components

Select the volumes and click **Next**. The Create Multi-Volume Snapshot - Summary dialog box appears (Figure 6-12).

Figure 6-12: Create Multi-Volume Snapshot - Summary

If the configuration is satisfactory, click **Finish** to create the custom snapshot collection. Click **Back** to correct any information.

The resulting snapshots will appear in the far left panel under the volume names, identified by the date and time when the snapshots were created.

The actual name of a custom snapshot collection will be the name you specified in the General Settings dialog box (Figure 6-10), followed by the date and time when the snapshots were created. For example, a custom snapshot collection named `mailsnapcol` will appear as `mailsnapcol-2006-06-20-11:17:04.67`.

Displaying Snapshot Collection Status

You can display the snapshot collections created from a volume collection. You can also display the custom snapshot collections created from a multi-volume snapshot operation.


Snapshot collections appear under the volume collection name in the far left panel, identified by the timestamp when the snapshot collection was created. Individual snapshots appear under the volume names with the same timestamp.

To display the snapshot collections created from a volume collection, click:

Volume Collection > *collection_name* > Snapshots tab

The Volume Collection Snapshots window appears (Figure 6-13), displaying the snapshot collections for the volume collection.

Figure 6-13: Volume Collection Snapshots



The snapshot collection status can be the following:

- **complete** — Each volume in the collection still has a snapshot resulting from the snapshot operation.
- **incomplete** — A snapshot that was part of the collection was deleted, either automatically or manually.


In the Snapshots panel, expand a snapshot collection to view the snapshots in the collection. Select a snapshot to display related tasks in the Activities panel. For example, you can modify the volume's snapshot settings, clone a snapshot, set a snapshot online, change the permission for the snapshot (read-write or read-only), or delete the snapshot.

To display details about a snapshot collection, double-click a snapshot collection in the Snapshots panel or click:

```
Volume Collection > coll_name > timestamp
```

The Volume Collection - Snapshot Collection window appears (Figure 6-14), displaying status of the snapshot collection and its snapshots.

Figure 6-14: Volume Collection - Snapshot Collection



The Snapshot Collection Status panel shows the integrity of the snapshot collection, in addition to the modification status of the snapshot collection:

- not modified — No snapshot in the collection is currently set online with read-write permission, so the data is sure to represent the volume contents at the time the snapshot was created.
- potentially modified — One or more snapshots in the collection are currently set online with read-write permission.

The Snapshots panel shows detailed information about the snapshots in the collection. Double-click an individual snapshot to display the Snapshot Status window, which shows details about the snapshot. See *Displaying Snapshot Configuration and Status* on page 6-6 for more information about snapshot status.

A custom snapshot collection is not associated with a volume collection. Therefore, to display information about a custom snapshot collection, you must know the timestamp of the snapshots in the collection.


To display a custom snapshot collection, click:

Volumes > *volume_name* > *snapshot_timestamp*

The Snapshot Status window appears. In the General Snapshot Information panel, click View snapshot collection. The Custom Snapshot Collection window appears (Figure 6-15), displaying the status of the custom snapshot collection and the volume snapshots.

A custom snapshot collection is indicated by the text `temporary collection` in the Volume collection field.

Figure 6-15: Custom Snapshot Collection



Modifying a Snapshot Collection

You can modify the name and description for a snapshot collection created from a volume collection and also for a custom snapshot collection. The name and description appear in the Snapshot Collection windows (Figure 6-14 and Figure 6-15).

When you create snapshots of the volumes in volume collection, the default name for the snapshot collection is the volume collection name followed by the date and time when the snapshot collection was created.

The default name for a custom snapshot collection is the name you specified when performing the multi-volume snapshot operation followed by the date and time when the custom snapshot collection was created.

To modify a snapshot collection created from a volume collection, click:

```
Volume Collection > collection_name > snapshot_collection_timestamp
```

The Volume Collection - Snapshot Collection window appears (Figure 6-14).


To modify a custom snapshot collection, click:

```
Volumes > volume_name > snapshot_timestamp > View snapshot collection
```

The Custom Snapshot Collection window appears (Figure 6-15).

From either window, click `Modify snapshot collection` in the Activities panel. The Modify Snapshot Collection Settings window appears (Figure 6-16).

Figure 6-16: Modify Snapshot Collection Settings



Modify the snapshot collection name and description, as needed, and click `OK`.

Deleting a Snapshot Collection

Deleting a snapshot collection created from a volume collection or deleting a custom snapshot collection will delete the snapshots in the collection.

To delete a snapshot collection created from a volume collection, click:

```
Volume Collection > collection_name > snapshot_collection_timestamp
```

The Snapshot Collection window appears (Figure 6-14).

To delete a custom snapshot collection, click:

```
Volumes > volume_name > snapshot_timestamp > View snapshot collection
```

The Custom Snapshot Collection window appears (Figure 6-15).

In either window, click `Delete snapshot collection` in the Activities panel. Confirm that you want to delete the snapshot collection.

7 Basic Replication

Using two PS Series groups configured as replication partners, you can set up a simple, yet robust disaster recovery plan that protects against catastrophic failures by replicating volume data from one group to another.

Only group administrators can configure replication partners, but pool administrators can configure replication on volumes. For more information, see *Managing Group Administration Accounts* on page 3-9.

Basic volume replication topics include the following:

- *Introduction to Replication* on page 7-1
- *Steps for Replicating Volumes* on page 7-9
- *Managing Replication Partners* on page 7-11
- *Configuring a Volume or Volume Collection for Replication* on page 7-19
- *Disabling Replication on a Volume or Volume Collection* on page 7-24
- *Creating Replicas* on page 7-24
- *Scheduling Replication* on page 7-25
- *Displaying Replication Status* on page 7-26
- *Pausing and Resuming Replication* on page 7-41
- *Cancelling an In-Progress Replication* on page 7-41
- *Cloning an Inbound Replica* on page 7-42
- *Deleting Replicas and Replica Sets* on page 7-45

For information about advanced replication tasks, including manual transfer replication and recovering data from the secondary group, see Chapter 8, *Advanced Replication*.

Introduction to Replication

There are many disaster recovery solutions available today. A solution must protect against day-to-day mistakes, such as users erroneously deleting volumes, in addition to computer viruses and site disasters.

Some solutions are time-consuming and require you to back up data and manually transport the backups to a different location. Other solutions require expensive hardware and the ability to synchronously copy data across long distances, which can adversely affect application performance.

Replication technology enables you to copy volume data across groups, protecting the data from failures ranging from the destruction of a volume to a complete site disaster, with no impact on data availability or performance.

Similar to a snapshot, a **replica** represents the contents of a volume at a specific point in time. A **replica set** is the entire set of complete replicas for a volume.

A volume and its replica set are always located in different groups. By separating the groups geographically, volume data is protected against a complete site disaster.

To replicate across two PS Series groups, you must configure the groups as **replication partners**. The groups can be physically located in the same building or separated by some distance, depending on your needs. There must be adequate network bandwidth and full IP routing between the groups.

Each replication partner plays a specific role in the replication of a volume. The volume is located in the **primary group**, while the volume replicas are stored in the **secondary group** in space that the secondary group delegated to the primary group. Mutual authentication provides security between partners.

You can create replicas of an individual volume or all the volumes in a volume collection. When you create replicas of the volumes in a collection, the resulting set of replicas (one for each volume) is called a **replica collection**. In addition, you can set up schedules to automatically create replicas or replica collections on a regular basis.

The first replica of a volume is a complete transfer of the volume data from the primary group to the secondary group over the network. For subsequent replicas, only the data that changed since the previous replica is transferred to the secondary group.

If you need to transfer a large amount of volume data for the first replication, and your network link has insufficient speed or bandwidth, you can use **manual transfer replication**. This enables you to copy the volume data to external media and then load the data from the media to the replica set on the secondary group. After the first data transfer is complete, replication continues, as usual, over the network.

To recover volume data from replicas in the secondary group, you can **clone** an individual replica to create a new volume. You can also use **failback** functionality to temporarily

host a volume from the secondary group for backup purposes or if the primary group is unavailable due to failure or maintenance. Later, you can fail back to the primary group and return to the original replication configuration, in some cases, by replicating only the volume changes that occurred while the volume was hosted on the secondary group. You can also use failback functionality to permanently switch partner roles.

Replication Configurations

For volume replication to occur, two groups must be configured as partners. There are three options for how partners can be configured:

- **One-way replication.** One group replicates volumes to another group.
- **Reciprocal replication.** Two groups replicate volumes to each other.
- **Centralized replication.** Multiple groups replicate volumes to another group.

Figure 7-1 shows one-way replication. AGroup and BGroup are configured as replication partners, with AGroup replicating two volumes to BGroup, which has delegated space to AGroup. From the perspective of AGroup, replication of volumes AVol1 and AVol2 is **outbound**. From the perspective of BGroup, replication of these volumes is **inbound**.

Figure 7-1: One-Way Replication




Figure 7-2 figure shows reciprocal replication, in which both groups replicate to each other. In this configuration, AGroup has now delegated space to BGroup for storing its replica sets, and replication of volume BVol1 is outbound from BGroup and inbound to AGroup.



Figure 7-2: Reciprocal Replication

Figure 7-3 shows a centralized replication configuration, in which two groups are replicating to a third group. In this configuration, CGroup has delegated space to both AGroup and BGroup.

Figure 7-3: Centralized Replication

Understanding Replication Space Requirements

Volume replication requires space on the primary group, where the volume resides, and on the secondary group, where the volume's replica set is stored, as shown in Figure 7-1.

Replication space is used as follows:

- On the secondary group, when configuring the primary group as a replication partner, you must **delegate space** to the primary group for the purpose of storing replica sets from the primary group. There is no default value. The amount of space to delegate to the primary group depends on the primary group's replica reserve requirements.
- On the primary group, when configuring a volume for replication, you must specify the size of the **replica reserve**, which is where the volume's replica set will be stored on the secondary group. The replica reserve is consumed from the space that the secondary group has delegated to the primary group. The replica reserve can range from 105% to an unlimited percentage of the volume reserve. The default and recommended value is 200%.
- On the primary group, when configuring a volume for replication, you must specify the size of the **local replication reserve** (5% to 200% of the volume reserve; the default is 100%) on the primary group. The local replication reserve has two purposes: to track volume changes during replication and, optionally, to store the failback snapshot (the most recent complete replica). The failback snapshot may enable you to synchronize the primary and secondary groups upon failback by replicating only volume changes, but it requires more local replication reserve.

You can also enable an option to use (borrow) free pool space if the local replication reserve is not sufficient.

See *Sizing Delegated Space and the Replica Reserve* on page 7-5 and *Sizing the Local Replication Reserve* on page 7-7 for information about sizing replication space.

Sizing Delegated Space and the Replica Reserve

On the secondary group, when configuring the primary group as a replication partner, you delegate space to the primary group for storing all replica sets from the primary group. Then, on the primary group, when configuring a volume for replication, you specify the size of the replica reserve, which is the portion of delegated space that will store the volume's replica set.

To determine the amount of space to delegate to the primary group, calculate the replica reserve requirement for each primary group volume that will be replicated and combine the values. It is recommended that you delegate more space than is minimally needed, so you can manually increase a volume's replica reserve to complete a replication operation.

The replica reserve for a volume can range from 105% of the volume reserve to an unlimited percentage. The default and recommended value is 200%.

To determine the appropriate size for a volume's replica reserve, it can help to understand how the replica reserve is used:

- When the first replica completes, it consumes 100% of the volume reserve from the replica reserve.
- When the next replication occurs, the replica reserve must have enough space to store the replicated data; otherwise, the replication cannot complete. You will be notified and given the opportunity to increase the replica reserve if space is not sufficient to complete the replication.
- When the replica reserve is consumed by replicas, the oldest replicas are deleted. However, the volume data represented by the most recent complete replica is always retained in the replica reserve and cannot be deleted.

Therefore, specifying the minimum replica reserve of 105% enables you to store more than one replica only if there are relatively few volume changes between replication operations.

If you specify the recommended value of 200% for a volume without thin provisioning, you are guaranteed the ability to retain at least two replicas in the replica reserve, even if the volume changes completely between replications. For a thin-provisioned volume, you may be able to retain only the most recent complete replica.

Note: Because replica reserve is based on the volume reserve, space requirements will increase if the volume reserve increases (for example, if you increased the size of a volume or if the volume reserve for a thin-provisioned volume automatically increased since the last replication).

It can be difficult to size the replica reserve because of these factors:

- Volume change rate between consecutive replication operations. This can be difficult to measure because workloads vary over time. Although some applications make a consistent number of volume writes, others have a workload that changes daily. Also,

for volumes with mainly random writes, you may need more replication space than the actual volume changes suggest.

- Type of volume I/O. There may be no direct correspondence between changes to the base volume and the amount of replica reserve needed. Especially in cases of random volume writes, you may find that you need relatively more replica reserve than expected, based on the actual volume changes.
- Number of replicas you want to retain on the secondary group. A small replica reserve space limits the number of replicas you can retain.
- IT operations. Some disk operations can affect replication. For example, defragmenting a disk or performing a database reorganization on a volume that has already been replicated may require a large transfer of data for the next replication. However, subsequent replications may transfer data more efficiently than before the disk operation.

To determine the right amount of replica reserve, you may want to set a value and adjust it after monitoring replication activity. In general, you may want to be generous when initially sizing the replica reserve.

Sizing the Local Replication Reserve

The local replication reserve has two purposes: track volume changes during replication and, optionally, store the failback snapshot (the most recent complete replica). The local replication reserve can range from 5% to 200% of the volume reserve and is consumed from free space in the same pool as the base volume. The default value is 100% of the volume reserve.

When you enable replication on a volume, you must specify the local replication reserve for the volume and whether to retain the failback snapshot. The failback snapshot establishes the failback baseline for the volume (that is, the point in time at which the primary and secondary groups are synchronized) and may enable you to synchronize the groups upon failback by replicating only volume changes.

If volume writes occur during a replication, the original (unmodified) volume data is stored in the local replication reserve, progressively consuming the space. The longer replication takes, the more volume changes can potentially occur. For example, if the entire volume is re-written, 100% of the volume reserve will be needed for the local replication reserve.

If you *do not* keep the failback snapshot, once the replication completes, any space used in the local replication reserve is released, and the reserve is not used again until volume writes occur during the next replication.

However, if you keep the failback snapshot, the space is not released. When the next replication completes, the failback replica is replaced by the most recent replica.

For the portion of local replication reserve that tracks volume changes, the recommended value is 100%. If you want to retain the failback snapshot, add 100% to the tracking changes value to obtain the recommended local replication reserve value.

For example, if you determine that you need a value of 100% for tracking changes, and you also want to retain the failback snapshot, the recommended value for the local replication reserve is 200%.

Note: Because local replication reserve is based on the volume reserve, space requirements will increase if the volume reserve increases (for example, if you increased the size of a volume or if the volume reserve for a thin-provisioned volume automatically increased since the last replication).

Local replication reserve usage depends on the volume changes during a replication, which can vary for each replication and are based on the workload and the time it takes for the replica data transfer to complete. For example, if no volume writes occur during a replication, no local replication reserve is used for tracking changes.

However, there may be no direct correspondence between changes to the volume and the amount of local replication reserve needed. Especially in cases of random volume writes, you may find that you need relatively more local replication reserve than expected, based on the actual volume changes.

Because the local replication reserve requirement can vary for each replication, you may want to use the recommended value, monitor the space usage, and adjust the reserve as needed.

If you specify less than the recommended value for the local replication reserve, it is recommended that you enable the feature that allows a volume to temporarily borrow free pool space if the size of the local replication reserve is not adequate. Once the data transfer is complete, any used pool space is freed, unless it is being used for failback.

If the local replication reserve is insufficient, there are several scenarios that may occur:

- If there is insufficient local replication reserve and the “borrow free pool space” feature is enabled (and free pool space is available), the replication will continue and an informational message will be logged, specifying that free pool space is being temporarily used.
- If there is insufficient local replication reserve and the “borrow free pool space” feature is enabled but there is *not* enough free pool space, the replication will be cancelled, data transfer will stop, and an error message will be logged, specifying that the replication was cancelled.
- If there is insufficient local replication reserve and the “borrow free pool space” feature is *not* enabled, the replication will be cancelled, data transfer will stop, and an error message will be logged, specifying that the replication was cancelled.
- If you attempt to fail back to the primary group by replicating volume changes, and there is insufficient local replication reserve to store those changes, you will receive an event message specifying that you need to increase the space.

If you find you have insufficient local replication reserve, you may be able to resolve the problem one of these ways:

- Increase the local replication reserve value.
- Enable the feature that allows a volume to borrow free pool space when needed.
- If free pool space is limited and multiple replication operations are “borrowing” space, stagger the replications so they are not all competing for pool space at the same time.

Steps for Replicating Volumes

To replicate a volume from one group to another, follow these steps:

1. **Configure two groups as replication partners.** Log in to each group and configure the other group as a replication partner. On each group, you must specify the partner name, partner’s group IP address, and reciprocal passwords for mutual authentication. If a group will be storing replicas from a partner, you must delegate space to the partner. Be sure the partners have the required network connectivity and replication space.

See *Sizing Delegated Space and the Replica Reserve* on page 7-5 and *Configuring Replication Partners* on page 7-11 for more information.

2. **Configure replication on the volume.** Log in to the primary group and configure replication on the volume. Specify the following information for each volume:
 - Partner that will store the replicas (secondary group). Specify the same partner for all the volumes in a volume collection.
 - Replica reserve size on the secondary group. See *Sizing Delegated Space and the Replica Reserve* on page 7-5 for more information.
 - Local replication reserve size on the primary group and whether to enable the “borrow free space” option. See *Sizing the Local Replication Reserve* on page 7-7 for more information.
 - Whether the first replication will occur over the network or by using manual transfer replication.
 - Whether to keep the failback snapshot. You will need more local replication reserve if you keep the failback snapshot, than if you do not.

If you are replicating a volume collection, you must also configure replication on the collection. See *Configuring a Volume or Volume Collection for Replication* on page 7-19 for more information.

3. **Start the replication.** Either create a replica or set up a schedule to create replicas on a regular basis. You can replicate an individual volume or a volume collection. See *Creating Replicas* on page 7-24 and *Scheduling Replication* on page 7-25 for more information.

Note: If you are using manual transfer replication to transfer the data for the first replication, do not create a replica until you are ready to start the process. See *Displaying Replication Status* on page 7-26 for more information.

4. **Check that the replication completes.** Periodically, you should also check the replication space utilization and make modifications as needed.

See *Displaying Replica Collections* on page 7-31 and *Displaying the Replication Configuration for a Volume* on page 7-29 for status information.

After replicating a volume, you can recover data from replicas, if necessary. See *Recovering Data on the Secondary Group* on page 8-28.

Managing Replication Partners

The first step in setting up replication is to configure two groups as replication partners. A group can have multiple replication partners, but a volume (or volume collection) can replicate to only one partner at a time.

You must have group administrator permissions to configure replication partners. For more information, see *Managing Group Administration Accounts* on page 3-9.

Replication partner management topics include:

- *Replication Partner Requirements* on page 7-11
- *Configuring Replication Partners* on page 7-11
- *Modifying a Replication Partner* on page 7-17
- *Deleting a Replication Partner* on page 7-18

Replication Partner Requirements

Replication partners must meet the following requirements:

- Network connectivity, with adequate network bandwidth between the groups, in addition to full IP routing. A slow network link between the partners may result in long replication times. You may want to use manual transfer replication for the first replication if you have a slow network link and a large amount of data to transfer.
- Sufficient space for replication. Volume replication consumes free space in the primary group, where the volume resides, and in the secondary group, where the replicas are stored. You will need more replication space if you increase the size of a volume, or if the volume reserve for a thin-provisioned volume grows automatically. In addition, recovering data from replicas may require additional space on the secondary group.

See *Understanding Replication Space Requirements* on page 7-5 for more information about space requirements.

Configuring Replication Partners

To configure a replication partner, click:

Replication Partners > Configure partner

A series of dialog boxes appear. Specify the following information:

- **Group name and IP address.** Enter the partner's group name and group IP address. Be sure to specify the actual, current group name. If the partner group administrator changes the group name or the IP address at a later time, you *must* make the reciprocal change in the partner configuration.
- **Contact information.** Optionally, enter contact information for the partner.
- **Reciprocal passwords for mutual authentication.** Two passwords are required for mutual authentication between groups configured as replication partners. Each partner provides a password to the other partner, which validates the password. You must specify:
 - Password for the partner. Enter a password for the partner, or allow a password to be automatically generated. The partner group administrator must specify this exact password in the `Password obtained from the partner` field when configuring the current group as a partner.
 - Password obtained from the partner. Enter the password that the partner group administrator specified in the `Password for the partner` field when configuring the current group as a partner.
- **Delegated space.** If the group will be storing volume replicas from the partner, specify the amount of space to delegate to the partner (primary group) and the storage pool from which to take it. Delegated space is immediately consumed from free pool space. For more information on sizing delegated space, see *Sizing Delegated Space and the Replica Reserve* on page 7-5.

Note that communication between replication partners is not established until replication is *enabled* on a volume. Therefore, keep in mind the following issues:

- Free pool space will be negative when specifying delegated space until communication is established between the partners.
- You will not know if password or configuration problems exist until replication is enabled on a volume. If you receive a login error message, check the reciprocal passwords on the partners and make sure they agree.
- Space delegated to a partner or received from a partner will not be visible in the Group Manager GUI or CLI until replication is enabled on a volume.
- If you increase the size of a volume, or if the volume reserve for a thin-provisioned volume has grown since the last replication, the secondary group will not recognize the change until the next replication.

- Replication management changes made on the secondary group will not be reflected in the Group Manager GUI or CLI on the primary group until the next replication.

The following figures show examples of using the Group Manager GUI to configure two groups as replication partners. In the examples, AGroup is replicating volumes to BGroup. The figures alternate between AGroup and BGroup, showing each step in the process.

Step 1: Specifying the partner identification. In this step, you specify the partner's group name, IP address, and an optional description.

Figure 7-4: Step 1 on AGroup - Partner Identification






Figure 7-5: Step 1 on BGroup - Partner Identification



Step 2: **Specifying the contact information.** In this step, you optionally specify the name, e-mail address, and phone numbers for the partner administrator.

Figure 7-6: Step 2 on AGroup - Contact Information



Configure replication partner

Step 2 - Contact Information

Contact information:

Name: Bob chen

E-mail address: chen@email.com


Phone: 123-456-7890

Mobile:

Back Next Finish Cancel

Java Application Window

Figure 7-7: Step 2 on BGroup - Contact Information



Configure replication partner

Step 2 - Contact Information

Contact information:

Name: Mary Jones

E-mail address: jones@email.com

Phone: 111-222-4444

Mobile:

Back Next Finish Cancel

Java Application Window

Step 3: Setting the mutual authentication passwords. In this step, you specify the reciprocal passwords for the partners.

In Figure 7-8 and Figure 7-9, the reciprocal passwords are shown to illustrate how to enter the passwords when configuring the partners. In the actual dialog boxes, the passwords will not be displayed.

Figure 7-8: Step 3 on AGroup - Reciprocal Passwords






Figure 7-9: Step 3 on BGroup - Reciprocal Passwords



Step 4: Delegating space to the partner. In this step you specify how much space to delegate to the partner. If the delegated space will exceed the capacity of the group, the pool space table cell showing free space will turn red.

Figure 7-10: Step 4 on AGroup - No Delegated Space



Configure replication partner

Step 4 - Delegate Space

Optionally, you can store replicas from the partner. Enter the amount of pool space you want to delegate to the partner for storing replicas.

Delegated space: ☐ MB ☒ GB

Storage pool assignment (cannot be changed once set):


Storage pool: default

Storage pool default			
	Current	New	Change
Delegated space	0 MB	0 MB	0 MB
Delegated space in use	0 MB	0 MB	0 MB
Free delegated space	0 MB	0 MB	0 MB
Free pool space	5.25 TB	5.25 TB	0 MB

Back Next Finish Cancel

Java Application Window

Figure 7-11: Step 4 on BGroup - Delegated Space



Configure replication partner

Step 4 - Delegate Space

Optionally, you can store replicas from the partner. Enter the amount of pool space you want to delegate to the partner for storing replicas.

Delegated space: ☐ MB ☒ GB

Storage pool assignment (cannot be changed once set):

Storage pool: default

Storage pool default			
	Current	New	Change
Delegated space	0 MB	200.01 GB	200.01 GB
Delegated space in use	0 MB	0 MB	0 MB
Free delegated space	0 MB	200.01 GB	200.01 GB
Free pool space	5.25 TB	5.06 TB	-200.01 GB

Back Next Finish Cancel

Java Application Window

Step 5: Verifying the partner configuration. In this step, ensure that the partner configuration is correct before clicking **Finish**. To make changes, click **Back**.

Figure 7-12: Step 5 on AGroup - Summary




Figure 7-13: Step 5 on BGroup - Summary



Modifying a Replication Partner

You can modify the configuration of a replication partner. For example, you may need to delegate additional space to the partner or modify partner settings if the partner's group name or group IP address has changed.

To modify a replication partner, click:

Replication partners > *partner_name*

The Replication Partner Status window appears. When you click a task link in the Activities panel, dialog boxes will be displayed, similar to those displayed when configuring a partner, as described in *Configuring Replication Partners* on page 7-11.

You can do one or more of the following:

- **Modify partner settings.** Click `Modify settings` and change the partner name, IP address, description, and contact information, as needed. Click `OK` to apply the changes.
- **Modify the space delegated to the partner.** Click `Modify delegated space` and enter the amount of space to delegate to the partner. The Space Utilization table shows how pool space is currently being used and how much space will be used with the new amount of delegated space. If the new delegated space will exceed the capacity of the pool, the table cell showing free pool space will turn red.

You cannot decrease the space delegated to a partner below what is currently being used to store the partner's replicas.

You cannot change the delegated space pool.

- **Modify mutual authentication passwords.** Click `Modify passwords` and change the passwords. Click `OK` to apply the changes.

You may need to make the reciprocal change on the partner by logging in to the partner group and modifying the passwords.

Deleting a Replication Partner

You can delete a replication partner if you want to sever the replication relationship between the two groups.

Note: If you delete a replication partner, any inbound partner replicas stored in the group will be deleted and the delegated space will be freed. However, outbound replicas stored on the deleted partner will continue to exist and can be accessed by logging in to the partner. The next replication of the volume will fail or be paused.

Before deleting a partner, you must pause any replication that is inbound from the partner, as described in *Pausing and Resuming Replication* on page 7-41.

To delete a replication partner, click:

Replication Partners > *partner_name* > Delete partner

Click OK to confirm that you want to delete the partner.

Configuring a Volume or Volume Collection for Replication

Before you can replicate a volume, you must configure the volume for replication. If you are replicating a volume collection, you must also configure replication on the collection. Each volume in the collection must be configured to replicate to the same partner.

To configure a volume for replication (or modify the existing replication configuration for a volume), click:

Volumes > *volume_name* > Configure replication

The Configure Volume Replication - General dialog box appears (Figure 7-14).

Figure 7-14: Configure Volume Replication - General

Configure volume replication

Step 1 - General Settings

Replication partner: **group2**

Replica reserve on the partner: **200** % of volume reserve (min. 105%)

Replication partner group2

	Current	New	Change
Replica reserve	0 MB	400.02 GB	400.02 GB
Free delegated space	723.74 GB	323.72 GB	-400.02 GB

Local replication reserve: **100** % of volume reserve (5 - 200%)

☒ Allow temporary use of free pool space

Storage pool usermail

	Current	New	Change
Local replication reserve	0 MB	200.01 GB	200.01 GB
Free pool space	1.77 TB	1.58 TB	-200.01 GB

Back Next Finish Cancel

Java Application Window

Specify information in the following fields:

- **Replication partner** – The partner you select must have already delegated space to the current group. See *Managing Replication Partners* on page 7-11 for more information.
- **Replica reserve on the partner** – Specify the portion of delegated space in the secondary group that will be allocated to the volume's replica set. The default and recommended amount is 200% of the volume reserve; the minimum is 105%. See *Sizing Delegated Space and the Replica Reserve* on page 7-5 for information about sizing the replica reserve.

Note: The table in the Configure Volume Replication - General dialog box (Figure 7-14) shows the available space on the primary and secondary groups. It will show a negative number for free delegated space if replication has not yet been attempted. This is normal and occurs because the groups do not communicate until replication is attempted.


If you are modifying an existing replication configuration, you cannot decrease the amount of replica reserve below what is currently being used to store replicas.

- **Local replication reserve** – Specify the amount of pool space in the primary group that will be used to track volume changes during replication and, optionally, store the failback snapshot. Local replication reserve is always consumed from the pool to which the volume is assigned. The default value is 100%.

For the portion of local replication reserve that tracks volume changes, the recommended value is 100%. If you want to retain the failback snapshot, to obtain the recommended value, add 100% to the tracking changes value. For example, if you determine that you need a value of 100% for tracking changes, and you also want to retain the failback snapshot, the recommended value for the local replication reserve is 200%. See *Sizing the Local Replication Reserve* on page 7-7 for more information.

- **Allow temporary use of free pool space** – By default, this option is enabled. Leave this checkbox selected to temporarily use free pool space in the primary group if the local replication reserve is not large enough to hold all the volume changes during replication. It is recommended that you use this feature if you specify less than the recommended value for the local replication reserve.

Click **Next**. The Configure Volume Replication - Advanced dialog box appears (Figure 7-15).

Figure 7-15: Configure Volume Replication - Advanced

Specify information in the following fields:

- **Initial data transfer** – Choose how the first data transfer will occur, either **Automatic** (over the network connection between the partners—the default method) or **Manual** (using manual transfer replication and physical transport of the volume data on external media).

The manual replication option applies only to the first transfer for the volume. Use this method if you have a very slow network connection or a large amount of data to transfer. All subsequent replications will be over the network connection. See *Displaying Replication Status* on page 7-26 for more information.


Note: To use the features in the latest PS Series firmware, including manual transfer replication and failback, both partners must be running version V3.2 or higher, and both partners must not allow firmware downgrades. See *Preventing Member Firmware Downgrades* on page 3-40 for more information.

- **Keep failback snapshot** – Select whether you want to keep the failback snapshot for the volume. If you choose this option, the local replication reserve will retain the failback snapshot (most recent complete replica). The failback snapshot may enable you to synchronize the groups upon failback by replicating only volume changes.

You can select this option at any time for a volume. The failback baseline will be established at the time of the next replication. If you do not retain the failback snapshot (or if the failback snapshot no longer exists), you must replicate all the volume data in order to synchronize the primary and secondary groups.

Click **Next**. The Summary screen appears (Figure 7-16).


Figure 7-16: Configure Volume Replication - Summary



Click **Finish** to complete the volume replication configuration. To make changes, click **Back**.

The Create Replica dialog box appears (Figure 7-17).

Figure 7-17: Create Replica



You can choose to create a replica now or later. You can also replicate the volume as part of a volume collection, or set up a schedule for creating replicas.


Note: If you selected `Manual` for the initial data transfer, do not create a replica until you are ready to start the manual transfer process. See *Displaying Replication Status* on page 7-26 for more information.

To configure a volume collection for replication, click:

Volume Collection > `collection_name` > Configure replication

The Modify Replication Settings - Volume Collection dialog box appears (Figure 7-18).

Figure 7-18: Modify Replication Settings - Volume Collection



Specify the partner to which the collection will be replicated. The collection and all its volumes must be configured to replicate to the same partner.

If a volume is already configured for replication, the Replicated To column will contain a link named for the replication partner. If a volume is not configured for replication, the Replicated To column will contain the `not replicated` link. Click either link to display dialog boxes that enable you to modify or configure replication for the volume.

Disabling Replication on a Volume or Volume Collection

To stop a volume or volume collection from replicating, disable replication on the volume or volume collection. Disabling replication on a volume or volume collection does not delete existing volume replicas on the partner.

Note: If you disable replication on a volume or volume collection and later re-enable replication, the initial replica will be a complete transfer of the volume data. If you want to temporarily stop volume replication, pause the replication instead of disabling it. See *Pausing and Resuming Replication* on page 7-41.

To disable replication on a volume, click:

```
Volumes > volume_name > Disable volume replication
```

A confirmation box appears. Click **Yes**.

To disable replication on a volume collection, click:

```
Volume Collection > collection_name > Disable collection replication
```

A confirmation box appears. Click **Yes**.

Creating Replicas

You can create replicas of individual volumes or volume collections at the current time. Alternately, you can set up a schedule for the creation of replicas, as described in *Scheduling Replication* on page 7-25.

The set of replicas for a volume is called a replica set. If you replicate a volume collection, the resulting set of replicas is called a replica collection. All the replica collections for a specific volume collection are called a replica collection set. When complete, a replica collection contains one replica for each volume in the collection and represents the contents of the collection volumes at the time the replication operation was initiated.

Before you can create a replica, you must set up the replication partners and enable replication on the volume (and volume collection, if applicable).

To create a replica of a volume at the current time, click:

```
Volumes > volume_name > Create replica now
```

Confirm that you want to create a replica.

To display replicas and replication status, select the volume in the far left panel, and then click the **Replication** tab.

To replicate a volume collection, click:

```
Volume collection > collection_name > Create replica now
```

The resulting replica collection will appear in the far left panel under the volume collection name. Replica collections are identified by the date and time that the replication operation was initiated.

Note: If a volume is configured to use manual transfer replication, after creating the first replica, you must perform additional tasks on both the primary group and the secondary group. See *Displaying Replication Status* on page 7-26 for more information.

Scheduling Replication

Schedules enable you to create replicas of a volume or volume collection at a specific time in the future or create replicas on a regular basis.

When scheduling replication, be sure to consider the capacity of the network link between the group and the replication partners. If a replication is still in progress when the next replication of the volume is scheduled to start, the scheduled replication will not start, and a message will be logged on the primary group. When a replication is skipped, more data will probably need to be transferred in the next replication.

Note: If you have a slow network connection between partners, consider using manual transfer replication for the first replication of a volume. See *Displaying Replication Status* on page 7-26 for more information. If you are using manual transfer replication for a volume, do not create a replication schedule for the volume until after the first replication is complete.

Before you can create a replica schedule for a volume or volume collection, you must enable replication on the volume and on the collection, as described in *Configuring a Volume or Volume Collection for Replication* on page 7-19.

For detailed information about creating a replica schedule, see *Managing Schedules for Snapshots and Replicas* on page 5-41.

Displaying Replication Status

You can display information about replication as follows:

- *Understanding Replication Status* on page 7-26
- *Displaying the Replication Configuration for a Volume* on page 7-29
- *Displaying Replica Collections* on page 7-31
- *Displaying Replication Activity Between Partners* on page 7-32
- *Displaying Outbound Replication Status* on page 7-35
- *Displaying Inbound Replication Status* on page 7-38

Understanding Replication Status

The Group Manager GUI displays the status of individual replicas, in addition to replication operations. The status may depend on whether the view is from the primary group or the secondary group.

Each replica is identified by the date and time the replication operation started.

Table 7-1 shows the possible replica statuses. Note that, from the view of the primary group, a replica status can also be one of the replication operation statuses described in Table 7-2.

Table 7-1: Replica Status

Status	Description
complete	All volume data has been stored on the secondary partner.
incomplete or in-progress	Replication is not complete.
manual-in-progress	Manual transfer replication is in progress, and you may need to perform additional tasks to complete the process. (In this case, an action item appears in the Outstanding Alarms panel at the bottom of the GUI.)
paused	Replication was paused
ready	Volume is not currently being replicated.

Possible values for replication operation status, as seen from the primary group, are shown in Table 7-2.

Table 7-2: Replication Operation Status – Primary Group

Status	Description
auth-failure	Authentication between the partners failed. Note: Check the passwords configured on the partners.
cancelling	Administrator cancelled the replication operation.
completed	The last replication completed successfully.
disabled	Replication of the volume has been manually disabled.
enabled	Volume is enabled for replication.
failed	Volume replication has failed. Note: Check the event log to determine why the replication operation failed.
in-progress	Volume is currently being replicated to the partner.
manual-in-progress	Manual transfer is currently in progress.
paused	Replication of the volume has been manually paused.
partner-down	Volume replication cannot proceed because the partner cannot be reached. Note: Check the network link between partners.
partner-needs-upgrade	Replication has been paused automatically because the partner is running incompatible firmware and must be upgraded to the same firmware as the group. Replication will resume automatically when the partner is upgraded.
partner-paused-outbound	Partner administrator paused outbound replication.
partner-paused-inbound	Partner administrator paused inbound replication.
paused	Administrator paused replication to the partner.

Table 7-2: Replication Operation Status – Primary Group (Continued)

Status	Description
paused-max-snap-reached	Volume replication paused because the partner (secondary group) has reached the maximum number of replicas or snapshots for a group. Note: To continue the replication, you must delete replicas or snapshots. If a schedule is being used to create replicas, you may want to decrease the maximum number of replicas to keep.
paused-partner-allows-downgrades	Replication has been paused automatically because the partner group needs to disallow firmware downgrades before replication can resume.
paused-remote-reserve-low	Volume replication was paused because the administrator of the primary group must increase the replica reserve. An event message will display the amount to which the replica reserve must be increased in order for the replication to complete. Note: See <i>Configuring a Volume or Volume Collection for Replication</i> on page 7-19 for information about increasing the replica reserve.
paused-remote-resize-failed	Volume reserve was increased, but the corresponding replica reserve resize operation on the partner failed due to lack of delegated space. Because replica reserve is based on a percentage of volume reserve, when the volume reserve increases either manually or automatically (for thin provisioned volumes), the replica reserve requirement automatically increases. Note: To continue replication, the partner group's administrator must increase the delegated space, as described in <i>Modifying a Replication Partner</i> on page 7-17.
ready	Replication of the volume is not currently in progress.
waiting	Volume replication data transfer has not yet started because the maximum number of in-use replication iSCSI sessions has been reached.

Possible values for replication operation status, as seen from the secondary group, are shown in Table 7-3.

Table 7-3: Replication Status – Secondary Group

Status	Description
auth-failure	Authentication between the partners failed. Note: Check the passwords configured on the partners.
far-end-down	No communication with the primary group for 1 hour or more, possibly due to a network problem or a failure at the primary group.
incomplete	Replication is not complete.
ready	Volume is not currently being replicated.
stopped	Replication was paused.


Displaying the Replication Configuration for a Volume

To display the replication configuration and status for a volume, click:

Volumes > *volume_name* > Replication tab

The Volume Replication – Volume Replicas window appears (Figure 7-19).

Figure 7-19: Volume Replication – Volume Replicas



The Replication Summary panel shows the following information about the volume:

- Replication Settings – Partner that is storing the volume’s replicas, the size of the replica reserve on the partner, and the size of the local replication reserve on the local group.
- Status – Whether fallback replication is enabled, the timestamp of the fallback snapshot stored on the primary group (establishes the fallback baseline for the volume, which is the point in time at which the primary and secondary groups are synchronized), and the size of any pending data transfers.
- Replication Schedules – How many schedules exist, how many are running currently, and the time of the next scheduled replication.

The Remote Replicas panel shows the replicas in the volume’s replica set and the status of the individual replicas. Each replica is identified by the date and time the replication operation started. Select the replica set or a replica to display related tasks in the Activities panel.

In the Remote Replicas panel, click `Replication History` to show details about each replication operation, including when it started, 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.

Note: If replication operations are taking too long, be sure you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link may result in long replication times.

Displaying Replica Collections


When you replicate a volume collection, the resulting set of replicas is called a replica collection. All the replica collections for a specific volume collection are called a replica collection set. When complete, a replica collection contains one replica for each volume in the collection and represents the contents of the collection volumes at the time the replication operation was initiated.

To display the replica collections for a volume collection, click:

```
Volume collection > collection_name > Replicas tab
```

The Volume Collection Replicas window appears (Figure 7-20).

Figure 7-20: Volume Collection Replicas



The Remote Replicas panel shows the volume collection's replication configuration and the replica collection set for the volume collection. Expand a specific replica collection to view the individual replicas.

Select the replica collection set to display related tasks in the Activities panel. For example, you can disable replication to the partner, delete an individual replica collection, or delete the entire replica collection set.

Select a specific replica to display related tasks in the Activities panel. For example, you can delete a replica.

Displaying Replication Activity Between Partners

To display replication activity between a group and all its replication partners, click **Replication Partners** in the far left column. The Replication Partner Summary window appears (Figure 7-21).

Figure 7-21: Replication Partner Summary

The screenshot shows the 'group1 Group Manager' window. The top status bar indicates 'Account: gpadmin' and 'Logged in 4/5/07 8:13:24 AM'. The left sidebar shows a tree view with 'Group group1' expanded, containing 'Group Configuration', 'Group Monitoring', 'Events (58 new)', 'Storage Pools', 'Members', 'Volumes', 'Volume Collections', and 'Replication Partners'. The 'Replication Partners' panel is selected, showing 'Partner group2'. The main content area has three panels: 'Activities' (with links for 'Configure partner', 'Modify settings', 'Modify passwords', and 'Delete partner'), 'Delegated Space' (showing 'Total outbound space: 1.0 TB' and 'Total inbound space: 1.0 TB'), and 'Replication' (a table showing replication details for 'group2').

Partner	Direction	Status	Delegated space	Free space
group2	group1 → group2	enabled	1.0 TB	848.64 GB (82.9%)
group2	group1 ↔ group2	enabled	1.0 TB	1.0 TB (100.0%)

In the Delegated Space panel, outbound space refers to the space on all the replication partners that is delegated to the group and used to store the group's replicas. Inbound space refers to space in the current group that is delegated to all the replication partners and used to store partner replicas.

In the Replication panel, there are two entries for each replication partner. The arrows in the Direction column indicate the direction of the replication. The top entry describes outbound replication from the group to the partner, and the bottom entry describes inbound replication from the partner to the group.

The Replication panel also shows the amount of space the partners have delegated to each other and the amount of free delegated space.

Note: If free delegated space is low, the secondary group administrator may be able to increase the space delegated to the primary group, as described in *Modifying a Replication Partner* on page 7-17.

Select the partner in the Replication panel to display related tasks in the Activities panel. For example, you can modify the partner configuration, delete the partner, or pause replication between the partners.

To display details about replication between the group and a specific partner, double-click the partner name in the Replication panel. You can also click:

Replication partner > *partner_name* > Status tab

The Replication Partner Status window appears (Figure 7-22), displaying the partner configuration, inbound and outbound replication status, including the number of volumes or volume collections replicated, and any in-progress replication operations.

Figure 7-22: Replication Partner Status

The screenshot shows the 'Replication partner group2' window. The left sidebar contains a tree view with 'group2' selected under 'Replication Partners'. The main area is divided into several sections:

- Activities:** A list of actions for 'Partner group2', including 'Modify settings', 'Modify passwords', 'Delete partner', 'Pause outbound', 'Modify delegated space', 'Pause inbound', 'Replication to ...', 'Delete replica set', 'Delete replicas', 'Pause volume replication', and 'Cancel replica creation'.
- General Partner Information:**
 - Partner Status:** Group name: group2, Group IP address: 172.17.96.210, Replication to partner: configured, Replication from partner: configured.
 - Contact Information:** Contact name: Bob Chen, E-mail: chen@email.com, Phone: 123-456-7890, Mobile: (blank).
 - Description:** Secondary group
- Replication Status:**
 - Outbound replication to group2:** Outbound replication: enabled, Delegated space: 743.99 GB (61.4% free), Volumes replicated: 9, Collections replicated: 1.
 - Inbound replication from group2:** Inbound replication: enabled, Delegated space: 1.26 TB (79.1% free), Volumes replicated: 3, Collections replicated: 0.
- Replication Progress:**

Total replications in progress between partners: 1

Direction	Volume	Started	Transferred	Status
group1 → group2	vol5	4/10/07 10:21:02 AM	507 MB / 507 MB	in-progress

The bottom status bar shows '1 outstanding alarm (1 new)'.

To display additional information about replication between the group and the partner, click the Outbound and Inbound tabs. See *Displaying Outbound Replication Status* on page 7-35 and *Displaying Inbound Replication Status* on page 7-38.

Displaying Outbound Replication Status

There are two ways to display outbound replication and replicas. On the primary group, where the base volume resides, click:

Volumes > *volume_name* > Replication tab

The Volume Replication window appears (Figure 7-19), as described in *Displaying the Replication Configuration for a Volume* on page 7-29.

Alternately, to display outbound replication from the primary group to a partner (secondary group), click:

Replication partners > *partner_name* > Outbound tab

The Outbound Replication – Volume Replicas window appears (Figure 7-23).

Figure 7-23: Outbound Replication – Volume Replicas

The screenshot shows the Group1 Group Manager interface. The left sidebar contains a tree view with 'Replication Partners' expanded, showing 'group2' selected. The main window is titled 'Replication partner group2' and has tabs for 'Status', 'Outbound', and 'Inbound'. The 'Outbound' tab is active.

Replica Space

Replica space on partner: 999.42 GB
 Used by replica sets: 662.61 GB (66.3%)
 Free delegated space: 336.81 GB (33.7%)

Utilization of replica space
 Used space 662.61 GB
 Free space 336.81 GB

Outbound replication to partner group2: **enabled**

Remote Replicas

☒ Volume replicas (8 of 12) ☐ Collection replicas (2 of 8) ☐ Replication history

Replicas	Details	Replication Status
Volume v1	1 replica. Reserved 40.02 GB	ready
Volume vol1	4 replicas. Reserved 600 MB	ready
Volume vol2	2 replicas. Reserved 1020 MB	ready
Volume vol3	1 replica. Reserved 810 MB	ready
Volume vol4	2 replicas. Reserved 20.01 GB	ready
Volume vol5	No replicas. Reserved 20.01 GB	manual-in-progress
Volume vol6	2 replicas. Reserved 400.02 GB	ready
Volume vol8	2 replicas. Reserved 120 MB	ready

At the bottom, there is a status bar showing '1 outstanding alarm (1 new)'.

The Replica Space panel shows the usage of the space that the partner has delegated to the current group.

Note: If there is not enough partner space delegated to the group, the partner administrator can increase the delegated space, as described in *Modifying a Replication Partner* on page 7-17.


In the Remote Replicas panel, click `Volume replicas` to display the volumes that the group is replicating to the partner. Replica status is described in Table 7-1. Replication status from the perspective of the primary group is described in Table 7-2.

Select a volume to display related tasks in the Activities panel. For example, you can pause or resume replication of the volume, cancel an in-progress replication, delete the replica set, or delete a specific replica.

Expand a volume to show the volume replicas (the replica set). Each replica is identified by the date and time the replication operation started. Replica status can be `complete` (all volume data has been stored on the secondary partner) or `incomplete` (the replication operation is in progress).

Note: Individual replicas are identified by the date and time that the replication operation was initiated. From the perspective of the primary group, the timestamp is based on the time zone set on the primary group. However, from the perspective of the secondary group, the timestamp is based on the time zone set on the secondary group.

In the Remote Replicas panel, click `Collection replicas` to display the volume collections that the group is replicating to the partner, as shown in Figure 7-24.

Figure 7-24: Outbound Replication – Collection Replicas

In the Remote Replicas panel, expand a volume collection to show the replica collections (the replica collection set). Expand a replica collection to show the individual replicas in the collection.

Select the volume collection name to display related tasks in the Activities panel. For example, you can disable replication of the volume collection, delete the entire set of replica collections, or delete a specific replica collection.

Select a replica to display related tasks in the Activities panel. For example, you can delete the replica from the replica collection.

In the Remote Replicas panel, click **Replication History** to show details about the last 10 replication operations for each volume (Figure 7-25), 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.

Note: If replication operations are taking too long, be sure you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link may result in long replication times.

Figure 7-25: Outbound Replication – Replication History

The screenshot shows the 'group1 Group Manager' interface. The left sidebar contains a tree view with 'Replication Partners' expanded, showing 'group2' selected. The main window is titled 'Replication partner group2' and has tabs for 'Status', 'Outbound', and 'Inbound'. The 'Outbound' tab is active. The 'Replica Space' section shows 'Replica space on partner: 1.0 TB', 'Used by replica sets: 175.37 GB (17.1%)', and 'Free delegated space: 848.64 GB (82.9%)'. A bar chart shows 'Used space' (blue) and 'Free space 848.64 GB' (green). Below this, it states 'Outbound replication to partner group2: enabled'. The 'Remote Replicas' section has three tabs: 'Volume replicas', 'Collection replicas', and 'Replication history' (selected). The 'Replication history' table lists the following data:

Started	Volume	Duration	Data size	Speed	Transfer
4/5/07 9:53:06 AM	vol2	26 min 15 sec	10.0 GB	1638 MB/min	complete
4/5/07 9:45:25 AM	exchangevol	2 min 25 sec	15.0 MB	839 MB/min	complete
4/5/07 9:44:55 AM	dbvol2	14 min 5 sec	2.0 GB	1238 MB/min	complete
4/5/07 9:38:53 AM	vol1	6 min 11 sec	20.0 MB	636 MB/min	complete
4/5/07 9:37:50 AM	exchangevol	7 min 10 sec	1.2 GB	533 MB/min	complete
4/5/07 9:37:20 AM	dbvol2	22 min 15 sec	9.0 GB	1638 MB/min	complete
4/5/07 12:00:25 AM	dbvol2	3 min 24 sec	15.2 MB	839 MB/min	complete
4/5/07 12:00:25 AM	exchangevol	14 min 5 sec	2.1 GB	1238 MB/min	complete
4/4/07 5:05:46 PM	dbvol2	6 min 13 sec	21.0 MB	636 MB/min	complete
4/4/07 5:05:01 PM	exchangevol	8 min 10 sec	4.2 GB	533 MB/min	complete

The bottom status bar shows 'No outstanding alarms'.


Displaying Inbound Replication Status

To display replication and replicas inbound from a partner to the group, click:

Replication partners > *partner_name* > Inbound tab

The Inbound Replication – Volume Replicas window appears (Figure 7-26).

Figure 7-26: Inbound Replication – Volume Replicas



The Delegated Space panel shows the usage of the space that the group has delegated to the replication partner. Also shown is failback replica space, which is the total amount of space for failback replica sets associated with volumes that were previously replicating to the partner.

Note: If the space that the group has delegated to the partner is low, consider increasing the delegated space, as described in *Modifying a Replication Partner* on page 7-17.

In the Local Replicas panel, click **Volume replicas** to display the volumes that the partner is replicating to the group, replication details, and status. Replica status is described in Table 7-1. Replication status from the perspective of the secondary group is described in Table 7-3.

Note: If replica reserve for a volume is low, on the primary group, consider increasing the replica reserve, as described in *Configuring a Volume or Volume Collection for Replication* on page 7-19.


Select a volume name in the Local Replicas panel to display related tasks in the Activities panel. For example, you can clone a replica, promote the replica set, delete a replica, or delete the replica set.

Expand a volume to show the volume replicas (the replica set). Select a replica to display related tasks in the Activities panel. For example, you can clone a replica or delete a replica.

Note: Individual replicas are identified by the date and time that the replication operation was initiated. From the perspective of the primary group, the timestamp is based on the timezone set on the primary group. However, from the perspective of the secondary group, the timestamp is based on the timezone set on the secondary group.

In the Local Replicas panel, click **Collection replicas** to display the volume collections that the partner is replicating to the group, as shown in Figure 7-27.

Figure 7-27: Inbound Replication – Collection Replicas



Select a volume collection to display related tasks in the Activities panel. For example, you can delete the replica collection set or an individual replica collection.

Expand a volume collection to show the replica collection set. Expand a replica collection to show the individual replicas.

Pausing and Resuming Replication

You must pause (temporarily stop) replication to perform certain tasks like promoting a replica set. You can later resume the replication if it does not resume automatically.

Table 7-4: Pausing Replication

To Pause:	Click:
Data transfer for a specific volume	Volumes > <i>volume_name</i> > Pause replication
All replication from a group to a partner (outbound)	Replication Partners > <i>partner_name</i> > Pause outbound
All replication from a partner to a group (inbound)	Replication Partners > <i>partner_name</i> > Pause inbound

Table 7-5: Resuming Replication

To Resume:	Click:
Data transfer for a specific volume	Volumes > <i>volume_name</i> > Resume replication
All replication from a group to a partner (outbound)	Replication Partners > <i>partner_name</i> > Resume outbound
All replication from a partner to a group (inbound)	Replication Partners > <i>partner_name</i> > Resume inbound

Cancelling an In-Progress Replication

You can cancel an in-progress volume replication.

Note: If you want to temporarily stop volume replication instead of cancelling it, you can pause the replication. See *Pausing and Resuming Replication* on page 7-41.

To cancel an in-progress volume replication, click:

Volumes > *volume_name* > Cancel replication

Cloning an Inbound Replica

You can clone an inbound replica to create a new volume in the secondary group with the same size and contents as the original volume at the time the replica was created.

Cloning a replica consumes additional space in the secondary group from the delegated space pool (equal to the volume reserve at the time the replica was created). In addition, if you want to create snapshots or replicate the new volume, additional space will be needed.

After cloning, the replica still exists and the replica set remains complete. In addition, cloning does not affect volume replication. The primary group, if it is still functioning, can still replicate the volume to the secondary group.

Use the following procedure to clone a replica:

1. Ensure that the secondary group has enough free pool space to clone the replica (100% of the volume reserve at the time the replica was created).
2. Log in to the secondary group and click:

```
Replication partners > partner_name > Inbound tab
```

The *partner_name* is the name of the primary group (the group where the original volume resides).

3. In the Local Replicas panel of the Inbound Replication window, click `Volume replicas`. Expand the volume and select the replica.

If the replica is part of a replica collection, click `Collection replicas` in the Local Replicas panel. Expand the volume collection and the replica collection. Then, select the replica.

4. In the Activities panel, click `Clone replica`.
5. In the Clone Volume Replica - General Settings dialog box (Figure 7-28), specify the new volume name and, optionally, a description. Then click `Next`.

Figure 7-28: Clone Volume Replica - General Settings

Clone volume replica 5/8/07 3:06:10 PM

Step 1 - Volume Settings

General properties:

* Volume name:

Description:

Storage pool assignment:

Storage Pool	Capacity	Free	Drives
☒ default	7.85 TB	5.57 TB	SATA

Back Next Finish Cancel

Java Application Window

- In the Clone Volume Replica - Space Reserve dialog box (Figure 7-29), change the default amount of snapshot reserve, as needed. (To change the volume size or thin provisioning setting, you must finish creating the volume and then modify the volume). Then, click **Next**.

Figure 7-29: Clone Volume Replica - Space Reserve

Clone volume replica 5/8/07 3:06:10 PM

Step 2 - Space reserve

Volume space:

* Volume size: ☐ MB ☒ GB (max: 5.57 TB)

☐ Thin provisioned volume (use sliders below to adjust settings)

Snapshot space:

* Snapshot reserve: % of volume reserve

Reported volume size 10.0 GB

Free 10.0 GB

Estimated changes in storage pool default


	Current	New	Change
Volume reserve	138.31 GB	148.32 GB	10.0 GB
Snapshot reserve	134.5 GB	144.51 GB	10.0 GB
Replication reserve	61.03 GB	61.03 GB	0 MB
Dedicated space	1.95 TB	1.95 TB	0 MB
Free pool space	5.57 TB	5.55 TB	-20.01 GB

Back Next Finish Cancel

Java Application Window


7. In the Clone Volume Replica - iSCSI Access window (Figure 7-30), specify access controls for the volume. Then, click **Next**.

Figure 7-30: Clone Volume Replica - iSCSI Access



8. In the Clone Volume Replica - Summary window (Figure 7-31), if all the information is correct, click **Finish**. To make changes, click **Back**.

Figure 7-31: Clone Volume Replica - Summary



The new volume will appear in the list of volumes in the far left panel.

Deleting Replicas and Replica Sets

You can delete replicas or replica sets stored on a partner (outbound replicas) or locally in the group (inbound replicas). You can also delete a replica collection stored on a partner.

See *Deleting Outbound Replicas* on page 7-45, *Deleting Inbound Replicas* on page 7-46, and *Deleting a Replica Collection* on page 7-47 for more information.

Deleting Outbound Replicas

You can delete replicas or replica sets stored on a partner (outbound replicas).

Notes: Deleting a replica set will stop the volume replication. You will have to re-enable replication to create replicas.


You cannot delete the most recent replica.

To delete an outbound replica or replica set, click:

Volumes > *volume_name* > Replication tab

The Volume Replication window appears (Figure 7-19). In the Remote Replicas panel, select the replica set and click *Delete replica set* or *Delete replicas*, as desired. If you are deleting a replica, the Delete Replicas dialog box appears (Figure 7-32). Select one or more replicas to delete.

Figure 7-32: Delete Replicas



Alternately, to delete an outbound replica set or replica, click:

```
Replication partners > partner_name > Outbound tab
```

Click `Volume replicas` to display the Outbound Replication - Volume Replicas window (Figure 7-23). In the Remote Replicas panel, select the volume. Then, click `Delete replica set` or click `Delete replicas` in the Activities panel. If you are deleting a replica, the Delete Replicas dialog box appears (Figure 7-32). Select one or more replicas to delete.

To delete an outbound replica that is part of a replica collection, click:

```
Volume Collections > collection_name > Replicas tab
```

The Volume Collection Replicas window appears (Figure 7-20). In the Remote Replicas panel, expand the replica collection and select the replica. In the Activities panel, click `Delete replica`.

Alternately, to delete an outbound replica that is part of a replica collection, click:

```
Replication partners > partner_name > Outbound tab
```

Click `Collection replicas` to display the Outbound Replication - Collection Replicas window (Figure 7-24). In the Remote Replicas panel, expand the volume collection and then expand the replica collection. Select a replica and, in the Activities panel, click `Delete replica`.

Deleting Inbound Replicas

You can delete partner replicas and replica sets that are stored locally in the group (inbound replicas). However, be sure that you have permission from the administrator of the group where the original volume resides before deleting inbound replicas.

Caution: Be careful when deleting an inbound replica or replica set, because this will stop the volume replication, delete the partner's replicas, and may impair the partner's ability to recover data.

Before deleting an inbound replica or replica set, you must pause any replication that is inbound from the partner, as described in *Pausing and Resuming Replication* on page 7-41.

Note: You cannot delete the most recent replica.

To delete an inbound replica set or replica, click:

```
Replication partners > partner_name > Inbound tab
```

In the Local Replicas panel, click `Volume replicas`. The Inbound Replication - Volume Replicas window appears (Figure 7-26). In the Local Replicas panel, select the volume. In the Activities panel, click `Delete replica set` or `Delete replica`. If you are deleting a replica, the Delete Replicas dialog box appears (Figure 7-32). Select one or more replicas to delete.

To delete an inbound replica that is part of a replica collection, click:

```
Replication partners > partner_name > Inbound tab
```

In the Local Replicas panel, click `Collection replicas`. The Inbound Replication - Collection Replicas window appears (Figure 7-27). Expand the volume collection and then expand the replica collection. Select a replica and, in the Activities panel, click `Delete replica`.

Deleting a Replica Collection

To delete a replica collection, on the primary group, click:

```
Volume collection > collection_name > Replicas tab
```

The Volume Collection Replicas window appears (Figure 7-20).

Select the replica collection in the Remote Replicas panel and click `Delete replica collection` in the Activities panel.

Note: You cannot delete a replica collection if it contains the most recent replica for a volume.

8 Advanced Replication

Using two PS Series groups configured as replication partners, you can set up a simple, yet robust disaster recovery plan that protects against catastrophic failures by replicating volume data from one group to another.

Only group administrators can configure replication partners, but pool administrators can configure replication on volumes. For more information, see *Managing Group Administration Accounts* on page 3-9.

Advanced volume replication topics include the following:

- *Performing Manual Transfer Replication* on page 8-1
- *Recovering Data on the Secondary Group* on page 8-28

For information about basic replication tasks, including setting up replication partners, configuring a volume for replication, and creating replicas, see Chapter 7, *Basic Replication*.

Performing Manual Transfer Replication

The first replication of a volume is a complete transfer of volume data from the primary group to the secondary group. Subsequent replication operations transfer only the volume changes since the last replication.

Normally, all data transfers occur over the network. However, if you have a lot of data to transfer for the first replication and a slow network link between the partners, you can use manual transfer replication to perform the first replication.

With manual transfer replication, you create a replica as usual. However, instead of the data transfer occurring immediately over the network, a temporary snapshot is created on the primary group, and a replica set and replica placeholder is created on the secondary group.

You then use the Manual Transfer Utility to connect to the temporary snapshot and copy the volume data to one or more transfer files. After the data is transferred to external media (for example, a CD-ROM), you use the Manual Transfer Utility to load the data into the replica set on the secondary group. Once you confirm that the manual transfer process is complete, the next replication will proceed as usual over the network.

There are two interfaces to the Manual Transfer Utility:

- Graphical user interface (GUI) that runs on Microsoft Windows through the Group Manager GUI. On Windows systems, the GUI is the recommended method of performing manual transfer replication.
- Interactive command line interface (CLI) that runs on Windows, Linux, and UNIX operating systems. If you attempted to use the GUI and received an error message, you must use the CLI. You can also use the CLI for scripting.

For installation information, see *Installing the Manual Transfer Utility* on page 8-2.

To use the GUI, see *Using the Manual Transfer Utility GUI* on page 8-3.

To use the CLI, see *Using the Manual Transfer Utility CLI* on page 8-19.

Note: Manual replication requires you to create an image copy of the volume for transport purposes. If you do not use the Manual Transfer Utility to perform manual transfer replication, you must use a raw data backup method. You cannot use a file-based backup method.

For UNIX operating systems, you can use a utility like `dd`. You can also use backup software that creates a mirror image of the volume. Multiple backup vendors provide software that can create image copies.

When using a raw data backup method to perform manual replication, follow the general instructions for using the CLI version of the Manual Transfer Utility. For example, you must create the correct access controls and confirm when the transfer is complete.

Installing the Manual Transfer Utility

To perform manual transfer replication, the Manual Transfer Utility (includes both the GUI and CLI versions) must be installed on the following:

- System (computer) with an iSCSI initiator and network access to the primary group to create the transfer files.
- System with an iSCSI initiator and network access to the secondary group to load the transfer files into the replica set.

You can obtain the Manual Transfer Utility from the Downloads area of the Technical Support website. Follow the installation instructions on the website.

Notes: If you are installing the Manual Transfer Utility on a UNIX or Linux operating system, the system must have Python 2.3 or later installed. You can download Python from the following website:

<http://www.python.org/>

If you want to encrypt the volume data, both systems (at the primary and secondary group locations) also must have pyCrypto installed. You can download pyCrypto from the following website:

<http://www.amk.ca/python/code/crypto>

Using the Manual Transfer Utility GUI

The Manual Transfer Utility GUI is the recommended method of performing manual transfer replication on a Windows system.

See *Installing the Manual Transfer Utility* on page 8-2 for installation information.

The general steps for using the GUI to perform manual transfer replication are as follows:

1. On a system with the Manual Transfer Utility installed and access to the primary group, perform the following steps:
 - a. Log in to the Group Manager GUI on the primary group and configure the volume for replication. Be sure to choose the manual transfer option for the first replication. See *Configuring a Volume or Volume Collection for Replication* on page 7-19.
 - b. Create a replica, as described in *Creating Replicas* on page 7-24. This will create a temporary snapshot on the primary group and a replica set and replica placeholder on the secondary group. In addition, action items will appear in the Group Manager GUI's Outstanding Alarms panel for both groups. The action items remain active until the manual replication is complete and committed on each group.

The temporary snapshot will be automatically deleted after the manual transfer operation completes.

- c. Start the Manual Transfer Utility GUI and create the transfer files on the system. See *Creating Transfer Files Using the GUI* on page 8-4.

- d. Copy the transfer files from the system to external media. For example, if you are using a CD-ROM, use your CD-ROM writer application to copy the files. If you created the transfer files directly on the external media, such as an external hard drive, skip this step.

If you divided the volume data among several transfer files, be sure to copy all the files for the volume onto the external media.

2. Transport the media to a system with access to the secondary group.
3. On a system with the Manual Transfer Utility installed and access to the secondary group, perform the following steps:
 - a. Make the transfer file data on the external media accessible to the system. For example, insert the CD-ROM or copy the files from tape to a directory.
 - b. Log in to the Group Manager GUI on the secondary group.
 - c. Start the Manual Transfer Utility GUI and load the data from the transfer files into the replica set. Then, confirm when you are done. This will update the replication status on both groups. See *Loading Transfer Files Using the GUI* on page 8-11.

The following sections describe some of these steps in detail.

Creating Transfer Files Using the GUI


Before you can create transfer files, you must install the Manual Transfer Utility on a system with access to the primary group, as described in *Installing the Manual Transfer Utility* on page 8-2. In addition, you must create the first replica.

Follow these steps to use the Manual Transfer Utility GUI to create transfer files:

1. From the system with the Manual Transfer Utility installed, log in to the Group Manager GUI on the primary group and click:

Volumes > *volume_name*


2. In the Activities panel, click `Perform manual transfer`. The Manual Transfer Utility GUI opens (Figure 8-1).

Figure 8-1: Manual Transfer Utility GUI – Primary Group

The top panel, entitled Manual Replications in Current Group, shows the manual transfer replication operations for the group to which the system is currently connected, including the direction of data transfer (inbound or outbound), volume name, replication status, replica timestamp, amount of data being transferred, and the status of the data transfer. The Actions column displays a task link.


The bottom panel, entitled Data Transfers on Local Machine, shows the data transfer operations from the system to the current group. If you select Show data transfers launched from other storage groups, data transfer operations from the system to other groups will appear.

3. In the Manual Replications in Current Group panel, select the volume for which you want to create transfer files and click the [Start data copy](#) link in the Actions column. The Copy Data to Transfer File - Source dialog box appears (Figure 8-2).

Figure 8-2: Copy Data to Transfer File – Source

4. By default, `Create access policies for local interfaces` is selected. Keep this checkbox selected to create temporary access control records that will allow the system's iSCSI initiator to connect to the temporary snapshot. If the checkbox is not selected, the initiator may not be able to connect to the volume snapshot, and manual transfer replication cannot proceed. The records are deleted automatically when the manual transfer operation is complete.

Click `Next`. The `Copy Data to Transfer File - Destination` dialog box appears (Figure 8-3).

Figure 8-3: Copy Data to Transfer File – Destination

5. In the Full path to transfer files directory field, specify either a drive letter (if you are writing the file directly to removable media) or full path where you want the transfer file to be created. (If you have already performed a manual transfer operation, the field will default to the directory used in the last operation.)

The transfer file name is created automatically and uses the following syntax:

```
primary_name.volume_name.secondary_name.n.eqd
```

If multiple transfer files are created for the same volume, *n* indicates the order in which each was created, starting from 0. For a single transfer file, *n* will be 0.


In addition, select the copy operation options for the data transfer file:

- Compress data to reduce size – This option is selected by default. You can turn off compression if any of the following is true:
 - You are writing the transfer file to a medium that has its own compression algorithm (such as a tape)
 - Data cannot be compressed (such as video or audio files)
 - There is sufficient space for the uncompressed transfer file

- Speed of the copy operation is more important than data compression
- **Encrypt data; encryption password** – Select this checkbox to apply a password to the transfer file and enter a password in the field. You will need to supply this password when loading the transfer files on the secondary group.
- **Transfer file size cannot exceed** – Select this checkbox to limit the size of each individual transfer file. Also, select the appropriate unit of measure. Use this option to specify the capacity of the media on which you will be sending the transfer files. By default, the utility limits the transfer file size to 80% of the destination directory or drive capacity.
- **Create multiple transfer segments** – If you chose to limit the transfer file size (as described above), optionally select this checkbox and enter the number of transfer files you want to create. If you do not know how many files will be needed, specify 100, which is the maximum. The utility will create only as many files as necessary, limited by the maximum size you specified.


Click **Next**. The Copy Data to Transfer File – Summary dialog box appears (Figure 8-4).

Figure 8-4: Copy Data to Transfer File – Summary




6. If the information is correct, click **Finish** to start copying the volume data to the transfer file. To make changes, click **Back**.
7. When the data transfer begins, the Manual Transfer Start confirmation box is displayed (Figure 8-5). Click **OK**.

Figure 8-5: Manual Transfer Start



8. Check the status of the copy operation in the Manual Transfer Utility GUI window (Figure 8-6). The Data transfer status column displays one of the following:
 - **Not started** – Transfer has not been started.
 - **Running** – Transfer is in progress. In the bottom panel in the GUI window, the **Progress** column displays the current percentage of completion.
 - **Need more files** – Additional transfer files need to be created. You must load blank media onto the same drive letter you specified previously (if applicable) and then click **Resume** (▶) in the **Actions** column.
 - **Interrupted** – Manual transfer process was stopped by the user. To continue creating transfer files, click **Resume** (▶). Note that if you stop and resume the process of copying data to the transfer files, the utility will create the next file in the sequence. The utility cannot resume copying data into an existing transfer file.
 - **Transfer is complete** – All necessary transfer files have been created for the volume.
 - **Error** – An error occurred.

Figure 8-6 shows the Manual Transfer Utility GUI window after all the transfer files have been created.

Figure 8-6: Manual Transfer Utility GUI – Transfer File Creation Complete

9. Optionally, when all the transfer files have been created, you can clear the action item in the Group Manager GUI by clicking `Commit status to group`.

The action item on the secondary group remains active until you have finished loading the transfer files into the replica set and committed that operation.

10. Exit the Manual Transfer Utility GUI.

The next step is to copy the files to external media and then transport the media to the system with access to the secondary group. You might need to use a separate application (such as a CD-ROM writer application) to copy the files to external media.

You can then load the transfer files into the replica set on the secondary group, as described in *Loading Transfer Files Using the GUI* on page 8-11.

Loading Transfer Files Using the GUI

To complete the process of manual transfer replication, use the Manual Transfer Utility GUI to load the data from the transfer files into the replica set on the secondary group and then confirm that the process is complete.

Before you can load data, you must install the Manual Transfer Utility on a system with access to the secondary group, as described in *Installing the Manual Transfer Utility* on page 8-2. In addition, make the transfer files on the external media accessible to the system. For example, insert the CD-ROM or copy the files from tape to a directory.

Follow these steps to load data into a replica set:

1. From the system with the Manual Transfer Utility installed, log into the Group Manager GUI on the secondary group and select the replica set by clicking:


```
Replication Partners > partner_name > Inbound tab > replica_set
```

The replica set is identified by the volume name.


2. In the Activities panel, click `Perform manual transfer`. The Manual Transfer Utility GUI opens (Figure 8-7).

The top panel, entitled Manual Replications in Current Group, shows the manual transfer replication operations for the group to which the system is currently connected, including the direction of data transfer, volume name, replication status, replica timestamp, amount of data being transferred, and the status of the data transfer. The Actions column displays a task link. The Data size column does not apply to the load operation and does not represent the size of the transfer file (or files).

The bottom panel, entitled Data Transfers on Local Machine, shows the data transfer operations from the system to the current group. If you select `Show data transfers launched from other storage groups`, data transfer operations from the system to other groups will appear.

Figure 8-7: Manual Transfer Utility GUI – Secondary Group


3. In the top panel, select the volume and click `Start data load`. The Load Data from Transfer File - Source dialog box appears (Figure 8-8).

Figure 8-8: Load Data from Transfer File - Source

4. Enter the path to the transfer file. For mounted removable media (for example, a CD-ROM), specify the drive letter.


If the data transfer file was encrypted, select `Encryption password` and enter the correct password.

Click `Next`. The Load Data from Transfer File - Destination dialog box appears (Figure 8-9).

Figure 8-9: Load Data from Transfer File – Destination

5. By default, the checkbox next to `Create access policies for local interfaces` is selected. Keep this checkbox selected to create temporary access control records that will allow the system's iSCSI initiator to connect to the replica set. If the checkbox is not selected, the initiator may not be able to connect to the replica set, and manual replication cannot proceed. The records are deleted automatically when the manual transfer operation is complete.

Click `Next`. The Load Data from Transfer File - Summary dialog box appears (Figure 8-10).

Figure 8-10: Load Data from Transfer File – Summary


6. If the information is correct, click **Finish** to start the data load operation. To make changes, click **Back**.
7. When the data load operation begins, the Manual Transfer Start confirmation box is displayed (Figure 8-5). Click **OK**.
8. Check the status of the data load in the lower panel of the Manual Transfer Utility GUI window (Figure 8-7).

If all the transfer files are available to the system, they are all loaded automatically.

If there are more files to load, **More files needed** will appear in the Data transfer status column. Make them available to the system and then click the **Resume** icon (▶) in the Actions column.


9. When the Manual Transfer Utility GUI shows the load operation status **done**, click **Commit status to group** in the top panel.

Figure 8-11 shows a completed manual transfer that is ready to be committed. The bottom panel displays the history of this transfer operation, including the copy step (on the primary group) and the load step (on the secondary group).

Figure 8-11: Manual Transfer Utility GUI – Load Complete

10. In the Commit Complete Data Transfer confirmation box (Figure 8-12), click **Yes** to confirm that all the data transfer files for the volume have been loaded into the replica set on the secondary group.

This sends a confirmation message to the primary group, and clears the manual transfer replication action item for this operation in the Outstanding Alarms panel on both groups. In addition, the operation is cleared from the top panel of the Manual Transfer Utility GUI window.

Figure 8-12: Commit Complete Data Transfer

Managing Completed Operations Using the GUI


You can check the status of a manual transfer operation from a system with the Manual Transfer Utility installed and access to the group. You can also commit data transfers or clear a completed data transfer from the Manual Transfer Utility GUI.

Log into the Group Manager GUI, expand the Tools panel, and start the Manual Transfer Utility GUI. The bottom panel of the Manual Transfer Utility GUI (Figure 8-11) shows the history of manual transfer operations performed on the system connected to the group.


To confirm that a data transfer is complete, click `Commit` status of complete data transfers at the top of the window.

If there are no uncommitted transfers, the Data Transfers Already Committed confirmation box appears (Figure 8-13). Click `OK`.

Figure 8-13: Data Transfers Already Committed




To see the details for a specific manual transfer operation, select the operation in the bottom panel and click the `Show details` icon () in the Actions column. The Data transfer details window appears (Figure 8-14).

Figure 8-14: Data Transfer Details for a Volume

To clear a completed data transfer from the Manual Transfer Utility GUI and, optionally, delete the transfer files associated with the operation, select the operation in the bottom panel and click the `Delete transfer task icon` (🗑️) in the Actions column.

The Delete Task confirmation box appears (Figure 8-15).

Figure 8-15: Delete Task

To delete the transfer files, select `Also delete related transfer files with volume data`.

Then, click `Yes`. The operation will disappear from the lower panel.

Using the Manual Transfer Utility CLI

If you attempt to perform manual transfer replication using the Manual Transfer Utility GUI and receive an error message that you cannot use the GUI, you can use the interactive command line interface (CLI) to the utility. The CLI is also referred to as `reptool`.

See *Installing the Manual Transfer Utility* on page 8-2 for installation information.

Note: Even if you use the CLI to create transfer files, you may be able to use the GUI to load the transfer files, as described in *Loading Transfer Files Using the GUI* on page 8-11.

If you plan to use the GUI to load the data on the secondary group, the transfer files must be named using the following format:

`primary_group.volume_name.secondary_group.n.eqd`

The variable *n* is a sequence number, starting from 0 (zero), for multiple transfer files for the volume.

General steps for using the CLI to perform manual transfer replication are as follows:

1. On a system with the Manual Transfer Utility installed and access to the primary group, perform the following steps:
 - a. Log in to the primary group and use the Group Manager GUI or CLI to configure the volume for replication. Be sure to choose the manual transfer option for the first replication. See *Configuring a Volume or Volume Collection for Replication* on page 7-19.
 - b. Use the Group Manager GUI or CLI to create a replica, as described in *Creating Replicas* on page 7-24.
 - c. Start the Manual Transfer Utility CLI and create the transfer files on the system.

The CLI creates only one transfer file at a time. Therefore, unless all the data will fit into one file, you must run the CLI several times to create all the transfer files needed for a volume.

See *Creating Transfer Files Using the CLI* on page 8-20 and *Creating Additional Transfer Files Using the CLI* on page 8-24.

- d. Copy the transfer files from the system to external media. For example, if you are using a CD-ROM, use your CD-ROM writer application to copy the files. If you

created the transfer files directly on the external media, such as an external hard drive, skip this step.

If you divided the volume data among several transfer files, be sure to copy all the files for the volume onto the external media.

2. Transport the media to a system with access to the secondary group.
3. On a system with the Manual Transfer Utility installed and access to the secondary group, perform the following steps:
 - a. Make the transfer file data on the external media accessible to the system. For example, insert the CD-ROM or copy the files from tape to a directory.
 - b. Log in to secondary group and use the Group Manager CLI to create an access control record that will allow the system access to the replica set. You cannot use the Group Manager GUI to perform this task.
 - c. Connect the system to the replica set.
 - d. Start the Manual Transfer Utility CLI, load the data from the transfer files into the replica set, and confirm when you are done. This will update the replication status on both groups. See *Loading Transfer Files Using the CLI* on page 8-25.

Subsequent replication operations will proceed as usual over the network.

The following sections describe some of these steps in detail.

Creating Transfer Files Using the CLI

Before you can create transfer files, you must install the Manual Transfer Utility on a system with access to the primary group, as described in *Installing the Manual Transfer Utility* on page 8-2. In addition, you must create the first replica. This will create a temporary snapshot (called a replica snapshot in the CLI) in the primary group and a replica set and replica placeholder in the secondary group.

Follow these steps to use the Manual Transfer Utility CLI to create transfer files:

1. On the system with the Manual Transfer Utility installed, use an iSCSI initiator to connect to the temporary snapshot (replica snapshot) on the primary group.

A snapshot has an iSCSI target name based on the volume name and timestamp. For example:

```
iqn.2001-05.com.equallogic:0-8a0906-43cfc1902-0aeb2fb809c46168-  
vol5-2007-04-06-16:59:40.25
```

Note: If the initiator cannot connect to the snapshot, check the snapshot's access control records and create one that allows the initiator to connect. See *Managing Access Controls for Volumes and Snapshots* on page 5-23 for information on creating access control records for a snapshot.

2. Start the Manual Transfer Utility CLI by running the `reptool` executable or, from the Windows Start menu, click:

```
All Programs > EqualLogic > Manual Transfer Utility
```

An example of running the CLI is shown after these steps.

3. When prompted by the CLI for the operation to perform, choose to create a transfer file (enter 1).
4. When prompted, specify information about the create transfer file operation:
 - On Windows systems, enter the number associated with the connected snapshot.

On other operating systems, specify the device name corresponding to the connected snapshot (for example, `/dev/sde`).

- Enter a name for the transfer file, specifying the full path, or drive letter if you are writing directly to removable media. The file name will be automatically appended with the extension `.eqd`.

If you plan to use the GUI on the secondary group to load the transfer file data to the replica set, you must use the following naming syntax:

```
primary_group.volume_name.secondary_group.n.eqd
```

For example (including the full path name):

```
C:\temp\group1.emailvol.group2.0.eqd
```

- Specify a limit on the size of the transfer file, if applicable. For example, you might need to divide the volume data into several files that span multiple media. The number shown in brackets indicates the amount of free space on the drive or in the directory you entered.

- Optionally, enter an encryption password. You will need this password when loading the transfer files into the replica set on the secondary group.
 - Choose whether or not to compress the volume data. Compression is enabled by default. You can disable compression if you are writing the transfer file to a medium that has its own compression algorithm (such as a tape), or if the speed of the copy operation is more important than data compression.
5. When the transfer file is finished, the CLI displays a message indicating whether or not more files are needed.

If you need to create more files, see *Creating Additional Transfer Files Using the CLI* on page 8-24.

If you are done creating files, exit the utility.

The following example shows a complete operation that created a single transfer file for the volume contents.

```
Welcome to reptool - the command line version of the Manual Transfer Utility

Copyright (C) 2006-2007 by EqualLogic, Inc.

Only use reptool if you attempted to use the Manual Transfer Utility
GUI and received an error message. See the Group Administration
manual for information about manual transfer replication.

Enter the number associated with the operation you want to perform.

1. Copy data from a replica snapshot to a transfer file.
2. Load a transfer file into a replica placeholder.
3. Collect diagnostic information for your support provider.

Operation: 1

Enter the number associated with the replica snapshot.

1. emailvol-2007-04-23-14:57:18.239

Replica snapshot: 1

Enter a name for the transfer file. You must specify the full path
name. It is recommended that you use the following file name format:

primary_group.volume_name.secondary_group.n.eqd
```


The variable `n` is a sequence number used if you will be creating multiple transfer files for the volume.

Transfer file name: `C:\temp\group1.emailvol.group2.0.eqd`

Optionally, enter the transfer file size limit in MB. The value in brackets shows the available space on the transfer file device.

Transfer file size limit in MB [44636]:

Enter the encryption password for the transfer file.

Press the ENTER or RETURN key if you do not want encryption.

Password:

Do you want to compress the transfer file (y/n) [y]: `n`

10244 MB of 10244 MB in, 0 MB of 44545 MB out, 0.0 MB/s

The copy operation is complete. No additional transfer files are needed.
Press the ENTER key to exit.

Note: In the CLI output, the input and output values probably will not be the same as the volume size. The transfer files will generally add up to less than the volume size, because unwritten (unused) blocks of the volume are not copied. Even if every block of the volume contains data, if compression is used, the transfer files will add up to less than the volume size.

After creating all the transfer files, the next step is to copy the files to external media (for example, a CD-ROM) and then transport the media to the system with access to the secondary group. You might need to use a separate application (such as a CD-ROM writer application) to copy the files to external media.

You can then load the transfer files into the replica set on the secondary group, as described in *Loading Transfer Files Using the CLI* on page 8-25.

Creating Additional Transfer Files Using the CLI

If volume data cannot fit into a single transfer file, you must create additional files. To create additional transfer files using the Manual Transfer Utility CLI, follow these steps on the *same* system on which you create the first transfer file:

1. Use an iSCSI initiator to connect the temporary snapshot (replica snapshot) on the primary group.
2. Start the Manual Transfer Utility CLI by running the `reptool` executable or, from the Windows Start menu, click:

All Programs > EqualLogic > Manual Transfer Utility

3. A message is displayed informing you that an earlier operation is incomplete and needs more files. Enter `y` at the prompt.
4. You will be prompted for information about the additional transfer file operation:

- Enter a name for the next transfer file. For each consecutive transfer file, you can specify a different drive or path name or mount empty media on the same drive letter you specified before.

If desired, you can use command line editing. For example, you can enter the first few characters of the previous transfer file name and press `Tab`. The prompt displays a list of matching file names. You can then edit the name; for example, to increment its sequence number. It is strongly recommended to use the same naming convention for all transfer files for a volume.

- Specify a limit on the size of the transfer file, if applicable. The number shown in brackets in the prompt indicates the amount of free space on the drive or in the directory you specified.
 - Optionally, enter an encryption password. You will need to supply this password when loading the data into the replica set on the secondary group.
 - Choose whether or not to compress the volume data. Compression is enabled by default. You can disable compression if you are writing the transfer file to a medium that has its own compression algorithm (such as a tape), or if the speed of the copy operation is more important than data compression.
5. When the transfer file creation is finished, exit the program.
 6. If the transfer needs still more files, return to step 2.

After creating all the transfer files, the next step is to copy the files to external media (for example, a CD-ROM) and then transport the media to the system with access to the secondary group. You might need to use a separate application (such as a CD-ROM writer application) to copy the files to external media.

You can then load the transfer files into the replica set on the secondary group, as described in *Loading Transfer Files Using the CLI* on page 8-25.

Loading Transfer Files Using the CLI

To complete the process of manual transfer replication, use the Manual Transfer Utility CLI to load the data from the transfer files into the replica set on the secondary group and then confirm that the process is complete.

Before you can load data, you must install the Manual Transfer Utility on a system with access to the secondary group, as described in *Installing the Manual Transfer Utility* on page 8-2. In addition, make the transfer files on the external media accessible to the system. For example, insert the CD-ROM or copy the files from tape to a directory.

Follow these steps to use the Manual Transfer Utility CLI to load data into a replica set:

1. From the system with the Manual Transfer Utility installed, log into the Group Manager CLI on the secondary group and use the following command to identify the replica set and replica placeholder for the volume:

```
partner select partner_name inbound-replicaset select  
replicaset_name show
```

The replication status will be `in-progress`. The replica set name is based on the volume name (for example, `emailvol.1`). The replica placeholder name is based on the replica set name and includes a timestamp.

2. Use the following Group Manager CLI command syntax to create an access control record that will allow the system to connect to the replica placeholder:

```
partner select partner_name inbound-replicaset select replicaset_name  
access create initiator init_name | ipaddress address
```

Specify the initiator's IP address or name. You can specify `0.0.0.0` for unrestricted access, although this is not recommended.

Note the access control record number that is created, so you can delete it when the transfer is complete.

3. Connect the system to the iSCSI target associated with the replica placeholder identified in Step 1.
4. Start the Manual Transfer Utility CLI by running the `reptool` executable or, from the Windows Start menu, click:

All Programs > EqualLogic > Manual Transfer Utility

5. When prompted by the CLI for the operation to perform, choose to load a transfer file (enter 2).
6. When prompted, specify information about the load transfer file operation:
 - a. On Windows systems, enter the number associated with the connected replica placeholder.

On other operating systems, specify the device name corresponding to the connected replica placeholder (for example, `/dev/sde`).

- b. Enter the name of the transfer file to load, including the path or drive letter.

You can use file name completion. For example, enter the first few characters of a directory or file name and press `Tab`. The prompt displays a list of matching names. You can then finish typing the name of the file you want to load.

7. When the load operation for the file completes, a message is displayed indicating if more transfer files are needed.

To load more transfer files for the same replica set, exit the utility and then restart it as described in Step 4. You will get a message saying that the transfer is incomplete. Specify that you want to continue the transfer and follow the prompts.

When all the transfer files for a volume are loaded, exit the program.

8. Disconnect the initiator from the replica set.
9. Use the Group Manager CLI to delete the access control record you created in Step 2.

To display the access control records for a replica set, enter the following command:

```
partner select partner_name inbound-replicaset select replicaset_name
access show
```

To delete a access control record for a replica set, enter the following command:

```
partner select partner_name inbound-replicaset select replicaset_name
access delete number
```

Do not delete the record that will support normal replication between the partners after the first manual transfer is finished (typically, record number 1, with a user name that matches the primary group name, and an authentication method of chap-local).

10. Confirm that the manual transfer replication is complete by using the following Group Manager CLI command format:

```
partner select partner_name inbound-replicaset select rs_name
first-manual-xfer done
```

This command sends confirmation to the primary group that the manual transfer has been completed. On both groups, the replication status will change to `ready` (for the replica set) and `complete` (for the first replica).

The following example shows a complete data load operation using the CLI.

```
Welcome to reptool - the command line version of the Manual Transfer
Utility
```

```
Copyright (C) 2006-2007 by EqualLogic, Inc.
```

```
Only use reptool if you attempted to use the Manual Transfer Utility
GUI and received an error message. See the Group Administration
manual for information about manual transfer replication.
```

```
Enter the number associated with the operation you want to perform.
```

1. Copy data from a replica snapshot to a transfer file.
2. Load a transfer file into a replica placeholder.
3. Collect diagnostic information for your support provider.

```
Operation: 2
```

```
Enter the number associated with the replica snapshot.
```

1. emailvol.1-2007-05-02-16:53:53.115

```
Replica snapshot: 1
```

```
Enter the name of the transfer file you want to load. You must
specify the full path name.
```

```
Transfer file name: e:\group1.emailvol.group2.0.eqd
```

```
0 MB of 0 MB in, 0 MB of 10244 MB out, 0.0 MB/s
```

```
The load operation is complete.
```

```
Press the ENTER key to exit.
```

Note: In the CLI output, the input and output values may not match the size of the original volume. This can occur if unused volume blocks were skipped when creating the transfer files, or if the data size was reduced by compression. When the load operation is complete, all of the volume that was in use will have been copied to the replica set.

Recovering Data on the Secondary Group

There are several methods you can use to recover volume data stored in replicas on the secondary group. These methods can be used to:

- Provide access to a volume from the secondary group and optionally fail back to the primary group.
- Back up volume data from the secondary group.
- Test your replication configuration.
- Switch roles in a replication configuration (that is, the original secondary group becomes the new primary group, and the original primary group becomes the new secondary group).

The methods of recovering data differ in results and in the amount of space required in the secondary group. For example, you can:

- **Clone a replica.** This creates a new volume on the secondary group with a new name, but the same size and contents as the original volume at the time the replica was created. The replica still exists after the cloning operation, and volume replication from the primary group is *not* affected. Cloning a replica requires additional space in the secondary group (equal to the volume reserve at the time the replica was created).

See *Cloning an Inbound Replica* on page 7-42 for more information.

- **Promote a replica set.** This transforms the replica set into a volume and a set of snapshots that correspond to each complete replica. The volume will contain the data represented by the most recent complete replica and can be accessed by initiators as usual. Delegated space is reduced by the size of the replica reserve; no additional space is needed on the secondary group.

There are two options for promoting a replica set:

- Temporary promotion. You can temporarily promote an inbound replica set to a recovery volume and snapshots, while retaining the ability to later demote the recovery volume and snapshots to the original replica set.

Use this option when the primary group is unavailable because of failure or maintenance or if you want to temporarily host the volume from the secondary group for backup purposes.

Any changes that are made to the recovery volume are tracked in the secondary group, so you can later fail back those changes if the original volume becomes available. If you chose to keep the failback snapshot when configuring replication on the volume, you may be able to fail back by replicating only the volume changes, instead of the entire volume contents.

- Permanent promotion. You can permanently promote a replica set to a volume that is *not* a recovery volume. In addition, if you promoted an inbound replica set to a recovery volume, you can make this promotion permanent. In both cases, a permanent promotion disables the ability to later demote the volume to the original inbound replica set.

Use this option if the original volume has been destroyed and you are sure you cannot recover the volume.

Procedures for Recovering Data

Information and steps for the procedures for recovering data are described in:

- *Steps for Hosting a Volume for Backup Purposes* on page 8-30
- *Steps for Failback* on page 8-31
- *Steps for Switching Roles in a Replication Configuration* on page 8-34

Steps for Hosting a Volume for Backup Purposes

You may want to temporarily host a volume from the secondary group for backup purposes. This enables you to use a third-party application to back up a stable copy of volume data while the original volume continues to be available to users.

The following procedure assumes that no writes will be made to the volume while it is hosted on the secondary group, so there is no need to replicate any changes to the primary group.

Steps for hosting a volume for backup purposes are shown next and also described in Table 8-1:

1. On the primary group, configure replication on the volume, choosing whether to retain the failback snapshot in the local replication reserve, and create a replica. Make sure the replication completes.
2. On the secondary group, temporarily promote the inbound replica set to a recovery volume. After the promote operation, the replica set will disappear from the list of inbound replica sets and the recovery volume will appear in the list of volumes.

See *Temporarily Promoting an Inbound Replica Set to a Recovery Volume* on page 8-36.

3. Use a backup application to back up the recovery volume.
4. When the backup is complete, on the secondary group, demote the recovery volume to an inbound replica set, returning it to its original state. The recovery volume will disappear from the list of volumes, and the replica set will reappear in the list of inbound replica sets.

See *Demoting a Recovery Volume to an Inbound Replica Set* on page 8-40.

5. On the primary group, replicate the volume to synchronize the groups.

Table 8-1 describes the backup steps on the primary and secondary groups.

Table 8-1: Backup Steps

Step	Primary Group	Secondary Group
1	Configure replication on the volume, choosing whether to keep the failback snapshot, and create a replica. Ensure that the replication completes.	
2		Promote the volume's inbound replica set to a recovery volume.
3		Backup the recovery volume using your backup application.
4		Demote the recovery volume to an inbound replica set, returning it to its original state.
5	Create a replica to synchronize the groups.	

Steps for Failback

You may want to temporarily host a volume from the secondary group and then fail back to the primary group, returning to the original replication configuration. For example, you may want to do this if the primary group is unavailable because of failure or maintenance or if the original volume is destroyed.

While the volume is hosted on the secondary group, the group will track volume changes. Later when the original volume becomes available and you are ready to fail back, you can replicate to synchronize volume data across the groups.

If the failback snapshot is available when you fail back to the primary group, you may be able to synchronize the groups by replicating only the volume changes, not the entire volume contents. The failback snapshot represents the failback baseline for the volume (that is, the point in time at which the primary and secondary groups are synchronized).

Steps for failback are shown next and also described in Table 8-2:

1. On the primary group, configure replication on the volume, choosing whether to retain the failback snapshot in the local replication reserve, and create a replica. Make sure the replication completes.

2. When the volume becomes unavailable (or is manually set offline), on the secondary group, temporarily promote the inbound replica set to a recovery volume. After the promote operation, the replica set will disappear from the list of inbound replica sets and the recovery volume will appear in the list of volumes.

The promote dialog boxes enable you to set up access controls for the recovery volume, if desired, so systems can connect to the recovery volume.

See *Temporarily Promoting an Inbound Replica Set to a Recovery Volume* on page 8-36.

3. When the original volume becomes available, on the primary group, demote the volume to a failback replica set. The volume will disappear from the list of volumes and reappear as an inbound replica set that is flagged as a failback replica set.

See *Demoting a Volume to a Failback Replica Set* on page 8-47.

4. To synchronize the groups, on the secondary group, configure the recovery volume to replicate to the primary group and create replicas as needed. As usual, you must allocate local replication reserve and replica reserve. However, you do not have to delegate space. You cannot use manual transfer replication for a recovery volume.

Note: If the failback snapshot does not exist in the primary group, the first replication will be a full volume transfer.

When you are ready to fail back to the primary group, set the recovery volume offline and create a final replica. Ensure that the replication completes.

5. On the secondary group, demote the recovery volume to an inbound replica set, returning it to its original state. The recovery volume will disappear from the list of volumes, and the replica set will reappear in the list of inbound replica sets.

See *Demoting a Recovery Volume to an Inbound Replica Set* on page 8-40.

6. On the primary group, promote the failback replica set to a volume, returning it to its original state. The failback replica set will disappear from the list of inbound replica sets and the volume will reappear in the list of volumes. If the option to keep the failback snapshot is enabled on the volume, a replica will be created automatically in order to re-establish the failback baseline. Because the volume data is synchronized across the groups, no data will be transferred.

See *Promoting a Failback Replica Set to a Volume* on page 8-48.

If, during the failback process, an unrecoverable disaster occurs on the primary group, you can make the inbound replica set promotion permanent, disabling the ability to demote the volume to the original inbound replica set. See *Making an Inbound Replica Set Promotion Permanent* on page 8-42 for more information.

Table 8-2 describes the failback steps on the primary and secondary groups.

Table 8-2: Failback Steps

Step	Primary Group	Secondary Group
1	Configure replication on the volume, choosing whether to keep the failback snapshot, and create a replica. Ensure that the replication completes.	
2		Promote the volume's inbound replica set to a recovery volume. Optionally, create access controls to enable systems to connect to the recovery volume.
3	When the original volume becomes available, demote the volume to a failback replica set.	
4		To synchronize the groups, configure the recovery volume to replicate to the partner and create a replica. If the failback snapshot does not exist on the partner, the first replica will require a complete data transfer. When you are ready to fail back, set the recovery volume offline and create the final replica. Ensure that the replication completes.
5		Demote the recovery volume to an inbound replica set, returning it to its original state.
6	Promote the failback replica set to a volume, returning it to its original state.	

Steps for Switching Roles in a Replication Configuration

You may want to switch roles in a replication configuration. For example, if GroupA was replicating a volume to GroupB, you may want to switch the configuration and replicate the same volume from GroupB to GroupA.

If the failback snapshot is available, you may be able to synchronize the groups by replicating only the volume changes, not the entire volume contents. The failback snapshot represents the failback baseline for the volume (that is, the point in time at which the primary and secondary groups are synchronized).

Steps for switching roles are shown next and also described in Table 8-3:

1. On the primary group, configure replication on the volume, choosing whether to retain the failback snapshot in the local replication reserve. When you are ready to start the transition, set the volume offline and create a replica. Make sure the replication completes.
2. On the secondary group, temporarily promote the inbound replica set to a recovery volume. After the promote operation, the replica set will disappear from the list of inbound replica sets and the recovery volume will appear in the list of volumes. The promote dialog boxes enable you to set up access controls, if desired, so systems can connect to the recovery volume.

See Temporarily Promoting an Inbound Replica Set to a Recovery Volume on page 8-36.

3. On the primary group, demote the volume to a failback replica set. The volume will disappear from the list of volumes and reappear as an inbound replica set that is flagged as a failback replica set.

See Demoting a Volume to a Failback Replica Set on page 8-47.

4. To synchronize the groups, on the secondary group, configure the recovery volume to replicate to the primary group and create replicas as needed. As usual, you must allocate local replication reserve and replica reserve. You cannot use manual transfer replication for a recovery volume.

Note: If the failback snapshot does not exist in the primary group, the first replication will be a full volume transfer.

5. When you are ready to complete the role change, on the secondary group, convert the recovery volume to a non-recovery volume, making the promotion permanent and disabling the ability to demote the volume to the original inbound replica set. (You can still demote the volume to a failback replica set.) The resulting volume will appear in the list of volumes and will no longer be flagged as a recovery volume.

See *Making an Inbound Replica Set Promotion Permanent* on page 8-42.

6. On the primary group, convert the failback replica set to an inbound replica set, making the demotion permanent and disabling the ability to promote the replica set to the original volume. (You can still promote the inbound replica set to a recovery volume.) The resulting inbound replica set will appear in the list of inbound replica sets and will no longer be flagged as a failback replica set.

See *Making a Volume Demotion Permanent* on page 8-51.

Table 8-2 lists the steps for switching roles on the primary and secondary groups.

Table 8-3: Steps for Switching Roles

Step	Primary Group	Secondary Group
1	Configure replication on the volume, choosing whether to keep the failback snapshot. When you are ready to start the transition to the secondary group, set the volume offline and create a replica. Ensure that the replication completes.	
2		Promote the volume's inbound replica set to a recovery volume. Optionally, create access controls to enable systems to connect to the recovery volume.
3	Demote the volume to a failback replica set.	
4		To synchronize the groups, configure the recovery volume to replicate to the partner and create replicas, as needed. If the failback snapshot does not exist on the partner, the first replica will require a complete data transfer.

Table 8-3: Steps for Switching Roles (Continued)

Step	Primary Group	Secondary Group
5		When you are ready to complete the role transition, make the inbound replica set promotion permanent, converting the recovery volume to a non-recovery volume and disabling the ability to demote the volume to the original inbound replica set.
6	Make the volume demotion permanent, converting the failback replica set to an inbound replica set and disabling the ability to promote the replica set to the original volume.	

Temporarily Promoting an Inbound Replica Set to a Recovery Volume

You can temporarily promote an inbound replica set to a recovery volume and snapshots, while retaining the ability to later demote the volume and return to the original replica set. Any changes that are made to the recovery volume are tracked in the secondary group, so you can later fail back those changes if the original volume becomes available.

Use this option if you want to temporarily host the volume from the secondary group because the primary group is unavailable or for backup purposes.

When you promote a replica set to a recovery volume, no additional space is required on the secondary group. Delegated space is reduced by the size of the replica reserve.

Initiators can connect to a recovery volume as usual. However, restrictions for recovery volumes include the following:

- Cannot change the volume size
- Cannot change thin provisioning setting
- Can only replicate to the original partner
- Cannot change the volume name
- Cannot change public alias
- Cannot change RAID preference
- Cannot use manual transfer replication

Follow these steps to promote a replica set to a recovery volume:


1. Log in to the secondary group and click:

Replication partners > *partner_name* > Inbound tab

The *partner_name* is the name of the primary group.


2. In the Local Replicas panel of the Inbound Replication window, click **Volume replicas**. Select the volume and click **Promote to volume** in the Activities panel.
3. In the Pause Inbound Replication Confirmation box (Figure 8-16), confirm that you want to pause inbound replication from the partner. Replication will resume automatically (except for the affected volume) after the replica set is promoted.

Figure 8-16: Pause Inbound Replication Confirmation




4. In the Promote Replica Set – Volume Options dialog box (Figure 8-17), choose the recovery volume options:
 - Choose whether to set the volume online.
 - Choose whether to retain the iSCSI target name of the original volume.
 - Select **Keep ability to demote to replica set**. If you do not select this option, the promotion will be permanent and you will not be able to demote the resulting volume to the original replica set. This option is selected automatically if you chose to keep the failback snapshot when configuring the volume for replication, and it cannot be de-selected.


Then, click **Next**.

Figure 8-17: Promote Replica Set – Volume Options

5. In the Promote Replica Set – iSCSI Access dialog box (Figure 8-18), specify the access controls for the volume. Then, click **Next**.

Figure 8-18: Promote Replica Set – iSCSI Access

6. In the Promote Replica Set – Summary dialog box (Figure 8-19), ensure that the information is correct and click **Finish**. To make changes, click **Back**.

Figure 8-19: Promote Replica Set – Summary

After the promotion completes, the replica set will disappear from the list of inbound replica sets and the recovery volume will appear in the list of volumes.

At this point, you can access the recovery volume, as needed. Any writes to the recovery volume are tracked. After you are finished using the recovery volume, you can demote the recovery volume to return to the original replica set, as described in *Demoting a Recovery Volume to an Inbound Replica Set* on page 8-40.

If changes are made to the recovery volume, and the original volume becomes available on the primary group, you can demote the original volume to a failback replica set and replicate the changes to synchronize the groups. If the failback snapshot exists on the primary group, only the volume changes need to be replicated. Otherwise, the first replication will transfer the entire volume contents. You can then demote the recovery volume to the original replica set and promote the failback replica set to the original volume to return to the original replication configuration. See *Demoting a Volume to a Failback Replica Set* on page 8-47 and *Promoting a Failback Replica Set to a Volume* on page 8-48 for more information.

If you are certain that the original volume is destroyed, you can make the promotion permanent, as described in *Making an Inbound Replica Set Promotion Permanent* on page 8-42. Once you do this, the volume cannot be demoted to the original inbound replica set.

Demoting a Recovery Volume to an Inbound Replica Set

To return a recovery volume to the original inbound replica set, you can demote the recovery volume.


If changes were made to the recovery volume, before demoting the recovery volume, you can replicate the changes to the primary group and synchronize the groups. If the failback snapshot exists on the primary group, only the volume changes need to be replicated. Otherwise, the first replication will transfer the entire volume contents.

You can replicate any changes before demoting the recovery volume, or you can choose the option to automatically fail back the changes before the demotion. You will be given the opportunity to configure replication on the volume if it is not already configured. Before you can replicate changes, the original volume must be demoted to a failback replica set, as described in *Demoting a Volume to a Failback Replica Set* on page 8-47.

To demote a recovery volume to an inbound replica set, follow these steps:

1. On the secondary group, select the recovery volume in the far left panel.
2. In the Activities panel of the Volume Status window, select one of the following tasks:
 - Failback to primary – Replicates the recovery volume changes to the primary group and then demotes the recovery volume.
 - Demote to replica set – Demotes the recovery volume, discarding any recovery volume changes.
3. In the Demote Recovery Volume dialog box (Figure 8-20), click **Yes** to confirm that you want to demote the volume.

Figure 8-20: Demote Recovery Volume




If you chose to fail back the changes, and the recovery volume has changed and is configured for replication, a replica will be created on the primary group. This will take

some time, depending on the data that must be transferred. Then, the recovery volume will be demoted to an inbound replica set.

If you chose to fail back the changes, and the recovery volume has changed but is *not* configured for replication, a message appears asking if you want to configure replication before continuing (Figure 8-21).

Figure 8-21: Demote Recovery Volume - Configure Replication




Click **Yes** to display the Configure Volume Replication dialog boxes, as described in *Configuring a Volume or Volume Collection for Replication* on page 7-19. Once you configure replication, a replica will be created automatically. Then, the recovery volume will be demoted to an inbound replica set.

Click **No** to skip configuring the volume for replication. The recovery volume will be demoted to a replica set and the volume changes will be lost.

If you chose not to fail back the changes, but the recovery volume was changed, a message appears asking if you want to discard the changes (Figure 8-22).

Figure 8-22: Demote Recovery Volume - Discard Changes



Click **Yes** to discard the volume changes and demote the recovery volume.

Click **No** to cancel the demote operation. You can then configure replication on the volume and fail back the changes to the primary group.

When a recovery volume is demoted, it will disappear from the list of volumes and the replica set will appear in the list of inbound replica sets.


Making an Inbound Replica Set Promotion Permanent

If you are certain that the original volume is destroyed on the primary group (or if you want to switch roles in a replication configuration), you can make the promotion of an inbound replica set permanent. This will convert the recovery volume to a volume that cannot be demoted to the original inbound replica set.


To make the promotion of an inbound replica set permanent, follow these steps:

1. On the secondary group, select the recovery volume in the far left panel.
2. In the Activities panel of the Volume Status window, click `Make promote permanent`.
3. In the Convert Recovery Volume - Volume Settings dialog box (Figure 8-23), specify a volume name, optional description, and storage pool. You *must* change the volume name. Then, click `Next`.


Figure 8-23: Convert Recovery Volume - Volume Settings



4. In the Convert Recovery Volume - iSCSI Access dialog box (Figure 8-24), specify access controls for the new volume.

Figure 8-24: Convert Recovery Volume - iSCSI Access

5. In the Convert Recovery Volume - Summary dialog box (Figure 8-25), ensure that the information is correct and click **Finish**. To make changes, click **Back**.

Figure 8-25: Convert Recovery Volume - Summary

The resulting volume will no longer be flagged as a recovery volume.

Permanently Promoting an Inbound Replica Set to a Volume

You can permanently promote an inbound replica set to a volume. The resulting volume cannot be demoted to the original inbound replica set. No additional space is required on the secondary group; delegated space is reduced by the size of the replica reserve.

Use this option if you are certain that the original volume has been destroyed. To replicate the resulting volume, you must configure it for replication, and the first replication of the volume will be a complete transfer of the volume data.

If you did not choose to retain the failback snapshot when configuring replication for a volume, you can permanently promote an inbound replica set in one step, as described below. If you chose to keep the failback snapshot, you must promote the inbound replica set to a recovery volume and then make the promotion permanent, as described in *Temporarily Promoting an Inbound Replica Set to a Recovery Volume* on page 8-36 and *Making an Inbound Replica Set Promotion Permanent* on page 8-42.

To permanently promote an inbound replica set for a volume with the keep failback snapshot option disabled, follow these steps:

1. On the secondary group, click:


```
Replication partners > partner_name > Inbound tab
```

The *partner_name* is the name of the primary group.

2. In the Local Replicas panel of the Inbound Replication window, click `Volume replicas`. Select the replica set and click `Promote to volume` in the Activities panel.
3. In the Pause Inbound Replication Confirmation box (Figure 8-16), confirm that you want to temporarily pause inbound replication from the partner. Replication will resume automatically (except for the affected volume) after the replica set is promoted.
4. In the Promote Replica Set – Volume Options dialog box (Figure 8-26), select the volume options:
 - Choose whether to set the volume online.
 - Choose whether to retain the iSCSI target name of the original volume.
 - De-select `Keep ability to demote to replica set`.


Then, click **Next**.

Figure 8-26: Promote Replica Set - Permanent - Volume Options




5. In the Promote Replica Set - Permanent Volume Settings dialog box (Figure 8-27), specify a name (can be the same as the original volume name), optional description, and storage pool for the volume. Then, click **Next**.

Figure 8-27: Promote Replica Set - Permanent - Volume Settings




6. In the Promote Replica Set - Permanent iSCSI Access dialog box (Figure 8-28), specify access controls for the volume. Then, click **Next**.

Figure 8-28: Promote Replica Set - Permanent - iSCSI Access



7. In the Promote Replica Set - Permanent Summary dialog box (Figure 8-29), ensure that the information is correct and click **Finish**. To make changes, click **Back**.

Figure 8-29: Promote Replica Set - Permanent - Summary



The inbound replica set will disappear and the new volume will appear in the list of volumes.


Demoting a Volume to a Failback Replica Set

After you promote a volume's inbound replica set to a recovery volume, if the original volume becomes available, you can demote the original volume to a failback replica set and replicate to synchronize volume data across the groups. If you kept the failback snapshot (and it is available), only the volume changes will be replicated. Otherwise, the first replication will be an entire volume transfer.

To demote a volume to a failback replica set, follow these steps:

1. On the primary group, select the volume in the far left panel, and click **Demote to replica set**.
2. In the Demote Volume Confirmation box, confirm that you want to demote the volume to a failback replica set. Note that the volume's inbound replica set on the secondary group must already be promoted to a recovery volume, as described in *Temporarily Promoting an Inbound Replica Set to a Recovery Volume* on page 8-36.

Figure 8-30: Demote Volume Confirmation



After the demotion, the volume will disappear from the list of volumes, and the failback replica set will appear in the Local Replicas panel of the Replication Partner Inbound window.

At this point, you can configure the recovery volume on the secondary group to replicate to the primary group. When you are ready to return to the original replication configuration, set the recovery volume offline, perform a final replication, demote the recovery volume to the original inbound replica set, as described in *Demoting a Recovery Volume to an Inbound Replica Set* on page 8-40, and then promote the failback replica set to the original volume, as described in *Promoting a Failback Replica Set to a Volume* on page 8-48.

If you are switching roles in the replication configuration, you can make the volume demotion permanent, converting the failback replica set to an inbound replica set. You will not be able to promote the inbound replica set to the original volume. See *Making a Volume Demotion Permanent* on page 8-51 for more information.

You cannot delete a failback replica set. Instead, make the demotion permanent and then delete the resulting inbound replica set.


Promoting a Failback Replica Set to a Volume

As the final step in failing back to the primary group, you must promote the failback replica set to a volume, returning it to its original state.


Note: Before promoting a failback replica set, be sure to replicate any changes from the recovery volume on the secondary group. Also, you must demote the recovery volume to the original inbound replica set. See *Demoting a Recovery Volume to an Inbound Replica Set* on page 8-40 for information.

To promote a failback replica set to the original volume, follow these steps:

1. On the primary group, select the partner in the far left panel, click the **Inbound** tab, select the failback replica set in the Local Replicas panel, and then click **Promote to volume**.
2. In the Pause Inbound Replication confirmation box, confirm that you want to pause inbound replication.
3. In the Promote Replica Set - Volume Options dialog box (Figure 8-31), select whether you want to set the volume online. Then, click **Next**.

Figure 8-31: Promote Replica Set - Volume Options

4. In the Promote Replica Set - Summary dialog box (Figure 8-32), if the information is correct, click **Finish**. To make changes, click **Back**.

Figure 8-32: Promote Replica Set - Summary

The failback replica set will disappear from the Replication Partner Inbound window and reappear in the list of volumes.

Note: After you promote a failback replica set to a volume, the volume is automatically replicated immediately to re-establish the failback snapshot (and the failback baseline). If the volume data is synchronized across the groups, no data will actually be transferred.

Changing the Pool for a Failback Replica Set


You can change the pool for a failback replica set. If you later promote the failback replica set to a volume, the volume will reside in the new pool.

Note: Before converting a failback replica set to an inbound replica set, you must change the pool in which the failback replica set resides if the current pool does not contain the space the group has delegated to the partner.

To change the pool for a failback replica set, follow these steps:

1. On the primary group, select the partner in the far left panel, select the **Inbound** tab, select the failback replica set in the Local Replicas panel, and click **Change storage pool** in the Activities panel.
2. In the Move Failback Replica Set dialog box (Figure 8-33), select the pool and click **OK**.

Figure 8-33: Move Failback Replica Set



Making a Volume Demotion Permanent

For failback purposes, you can demote a volume to a failback replica set, as described in *Demoting a Volume to a Failback Replica Set* on page 8-47. If you are switching roles in the replication configuration, you can convert a failback replica set to an inbound replica set and make the demotion permanent. This will disable the ability to promote the replica set to the original volume.


Notes: You must convert a failback replica set to an inbound replica set in order to delete it.

Before converting a failback replica set to an inbound replica set, you must change the pool in which the failback replica set resides if the current pool does not contain the space the group has delegated to the partner.

To make a volume demotion permanent, follow these steps:

1. On the primary group, select the partner in the far left panel, select the **Inbound** tab, select the failback replica set in the Local Replicas panel, and click **Convert to replica set**.
2. In the Convert to Replica Set confirmation box, confirm that you want to convert the failback replica set.

Figure 8-34: Convert to Replica Set



The replica set will continue to appear in the Replication Partner Inbound window, but it will no longer be flagged as a failback replica set.

9 Member Management

A PS Series group consists of one or more PS Series arrays connected to an Ethernet network and configured as group members.

Only group administrators can manage members. For more information, see *Managing Group Administration Accounts* on page 3-9.

Member management topics include the following:

- *Introduction to Member RAID Policies* on page 9-1
- *Setting the RAID Policy and Pool for a New Member* on page 9-3
- *Displaying Member Configuration and Status* on page 9-6
- *Managing the Disk Array* on page 9-12
- *Managing the Member Network Configuration* on page 9-18
- *Managing Control Modules* on page 9-23
- *Moving a Member Across Pools* on page 9-27
- *Modifying a Member Name and Description* on page 9-29
- *Modifying the Group Membership Password* on page 9-30
- *Removing a Member from a Group* on page 9-30
- *Using LEDs to Identify a Member* on page 9-31
- *Displaying Member Enclosure Status* on page 9-32
- *Obtaining Member Service Information* on page 9-36

Introduction to Member RAID Policies

Each group member is assigned a RAID policy. A RAID policy consists of the following:

- RAID level (RAID 10, RAID 50, or RAID 5)
- Whether or not spare disks are automatically configured and used

After adding an array to a group, the storage in the member will be available once you set the RAID policy. Then, the disks are automatically configured according to the designated RAID level. Spare disks are also configured automatically.

Note: You can convert a member from one RAID policy to another only if the new RAID policy requires *less* disk space than the current policy, so be careful when choosing a RAID policy.

Although all RAID policies provide good performance and high availability, there are some differences. When choosing a member RAID policy, you should first identify the performance and availability needs of your workload and select a RAID policy that meets those needs. In a multi-member pool, you may want to configure the pool members with different RAID levels to take advantage of automatic performance load balancing.

The following RAID policies can be set using the Group Manager GUI:

- **RAID-10.** Striping on top of multiple RAID 1 (mirrored) sets, with one or two spare disks. RAID-10 provides good performance for random writes, in addition to the highest availability. However, since the disks are mirrored, RAID 10 provides the least capacity.
- **RAID-50.** Striping on top of two RAID 5 (distributed-parity) sets, with one or two spare disks. RAID-50 provides a good balance of performance (especially for sequential writes), availability, and capacity.
- **RAID-5.** One RAID 5 set, with one spare disk. RAID-5 is similar to RAID-50, with more capacity (two additional disks) but lower availability and performance.

Using the Group Manager CLI, you can specify two additional RAID policies:

- **RAID-10, no-spares.** Striping on top of multiple RAID 1 sets, with no spares, if possible.
- **RAID-50, no-spares.** Striping on top of two RAID 5 sets, with no spares, if possible.

Note: A no-spares RAID policy leaves the member vulnerable to data loss and should be used only at installations where extra disks and personnel are available at all times to replace failed disks.

Once you select a RAID policy for a member you can perform only the following RAID policy conversions:


- RAID-10 to RAID-50
- RAID-10 to RAID-5
- RAID-50 to RAID-5
- RAID-10, no-spares to RAID-50, no-spares (must use the CLI to convert)

If a member has a RAID policy that cannot be converted, you can change the policy only by removing the member from the group, re-adding it to the group, and then setting the desired RAID policy.

Setting the RAID Policy and Pool for a New Member


When you first add a member to a group, the Group Summary window will indicate that a member exists with an unconfigured RAID policy, as shown in Figure 9-1. The unconfigured member will appear in the default storage pool.

Figure 9-1: Group Summary - Member RAID Policy Not Set




To set the RAID policy for a new member, follow these steps:

1. Either double-click the member in the Storage Pools and Group Members panel or expand **Members** in the far left panel and select the member name. A warning message appears, stating that the member is not configured (Figure 9-2).


Figure 9-2: Warning RAID Policy Not Configured

2. Click **Yes** to display the Configure Member - General Settings dialog box (Figure 9-3). Select the pool for the member.

By default, or if there is only one pool, the member will be assigned to the default pool. You can also create a new pool at this point by clicking the **Create new pool** link. See Chapter 4, *Storage Pool Management* for more information about pools.

Figure 9-3: Configure Member - General Settings


3. Click **Next** to display the Configure Member - RAID Configuration dialog box (Figure 9-4). Select the desired RAID policy. The table will be updated, showing the member capacity with the selected RAID policy. See *Introduction to Member RAID Policies* on page 9-1 for information about choosing a RAID policy.

Figure 9-4: Configure Member - RAID Configuration

In addition, under Expand group capacity, select one of the options:

- Immediately expand group capacity — To immediately start using the member's disk space, select this option (the default and recommended setting). Member performance will not be optimal while the RAID verification is in progress, but the group is fully functional. The member status shows as `initializing` briefly at first, but changes to `online` while initialization is completing. Once you choose this setting for a member, you cannot change it.
 - Wait until the member storage initialization completes — To delay using the member's disk space until the RAID verification completes and the batteries are fully charged, select this option. The member space will be zero and the member status will show as `initializing` until the RAID verification completes and the batteries are charged. You can set this option *only* when choosing the member RAID policy for the first time. However, you can cancel this option by modifying the member's RAID configuration at a later time.
4. Click `Next` to display the Configure Member - Summary dialog box (Figure 9-5). If the member configuration is satisfactory, click `Finish`. To make changes, click `Back`.


Figure 9-5: Configure Member - Summary



Displaying Member Configuration and Status

To display the overall status of the group members, click **Members** in the far left panel. The Member Summary window appears (Figure 9-6).

Figure 9-6: Member Summary



In the Group Disk Space panel, select the type of information you want to display:

- **Group space utilization** – Group space allocated to volumes, snapshots, local replication reserve, delegated space, and the amount of free group space.

Note: The group free space shown in the GUI may not reflect the actual free space because the GUI uses a rounding algorithm to calculate free space. To display the actual free space in megabytes, use the Group Manager CLI command `cli-settings displayinMB`. Then, use the `show` command to display the free group space.

If free group space is low, you can add disks to members, as described in *Installing and Removing Disks* on page 9-16, or add members to the group, as described in *Adding Members to an Existing Group* on page 2-5.

- **RAID level space distribution** – Group space configured with each RAID level (RAID 10, RAID 5, and RAID 50).
- **Storage pool capacity** – Percentage of group space in each storage pool and the capacity of each pool.
- **Free storage pool space** – Free space in each storage pool.

Note that if a pool's free space is low, you can increase free space by moving volumes from the low-space pool to a different pool, as described in *Moving a Volume Across Pools* on page 5-29. You can expand pool capacity by moving an existing member to the low-space pool, as described in *Moving a Member Across Pools* on page 9-27.

For each member, the Group Members panel shows the storage pool, capacity, free space, RAID policy, disk type, status, firmware version, and iSCSI connections. Select a member to display related tasks in the Activities panel.

The member status is shown in Table 9-1.

Table 9-1: Member Status

Status	Description
unconfigured	Member was added to the group but has no RAID policy configured yet.
initializing	Member is initializing its storage according to the selected RAID policy. This status remains in effect until initialization is complete if you selected the option to wait before displaying the available storage on the member. See <i>Setting the RAID Policy and Pool for a New Member</i> on page 9-3.

Table 9-1: Member Status (Continued)

Status	Description
online	Member is a functioning member of the group. However, a member may experience a failure but still function in the group. Check the member health status for more information.
online, moving	Member is moving to a different pool. The time to move a member depends on the amount of data that must be moved from the member to the remaining pool members.
offline	Member is unavailable, was shut down, or experienced a failure. Note: If a member fails, you must quickly identify and correct the problem.
vacating	Member is in the process of moving data to the remaining group members prior to being removed from the group. The time to remove a member depends on the amount of data that must be moved from the member to the remaining pool members.
vacated	Member has successfully moved its data to the remaining pool members prior to being removed from the group.

To display detailed information about a specific member, double-click a member in the Group Members panel or click:

Members > *member_name* > Status tab

The Member Status window appears. In the Member Status panel, click **Front view** to display the front panel of the array model. Click **Rear view** to display the back panel of the array model, including the control modules and the power supply and cooling modules.

An uninstalled or unconfigured hardware component is indicated by a red X in the array graphic. A failed or failing component will appear in the array graphic with a warning or error status symbol. Place the mouse over a component to show status details.

Figure 9-7 shows the front panel and Figure 9-8 shows the rear panel of a PS5000 array.

Figure 9-7: Member Status - PS5000 Array Front Panel





Figure 9-8: Member Status - PS5000 Array Back Panel



The General Member Information panel displays the following:

- Member status
- Control module model number
- Number and type of disks
- Number of iSCSi connections.
- Member name and pool
- RAID policy and status

The RAID status is shown in Table 9-2.

Table 9-2: RAID Status

Status	Description
OK	Disk array has successfully initialized and performance is normal.
expanding	Disk array is expanding (for example, because additional disks have been installed or the RAID policy is changing from RAID-10 to RAID-50).
verifying	Disk array is initializing (for example, when the member's RAID policy is set).
degraded	A RAID 1 set or a RAID 5 set is degraded (for example, because a disk has failed and there is no spare to replace it). If the RAID policy is RAID-5 or RAID-50, performance may be impaired. Note: If the RAID status is <code>degraded</code>, immediately identify and replace any failed disks. It is recommended that spare disks be installed and functioning in an array.
reconstructing	Disk array is reconstructing data on a disk (for example, because a disk has failed and a spare is being used to replace it). During reconstruction, performance may decline. After reconstruction, performance returns to normal, unless a RAID set is degraded. Note: If the RAID status is <code>reconstructing</code>, immediately identify and replace any failed disks. It is recommended that spare disks be installed and functioning in an array.
failed	Multiple disk failures in a RAID 1 set or a RAID 5 set have caused the disk array to fail. The member will be set offline. Note: If the RAID status is <code>failed</code>, you may need to recover data from a backup or replica.

Table 9-2: RAID Status (Continued)

Status	Description
catastrophic loss	Disk array has lost group metadata or user data. The array will not initialize and an error message will appear on the console at regular intervals. Note: See the <code>clearlostdata</code> command in the <i>CLI Reference</i> manual for information about catastrophic data loss. You may need to recover data from a backup or replica.

The Member Health Status panel displays the following:

- Links to display images of the front and rear of the array
- A table of alarm statuses for the array components

If the status is warning or critical, an alarm will occur in the Outstanding Alarms panel. See *Monitoring Alarms and Action Items* on page 10-4 for information about the health conditions that result in alarms. Be sure to investigate any alarms.

If failed disks exist, click the `Disks` tab to display the disk configuration. You must replace any failed disks as soon as possible. See *Handling Failed Disks* on page 9-17 for more information.

The Member Space panel displays the following:

- Total member capacity (does not include the space in spare disks because they are used only if a disk failure occurs).
- Utilization of group space, including space used by volumes, snapshots and replicas, and free space.

If free member space is low, you can increase the member capacity by adding disks (up to the maximum allowed number). See *Installing and Removing Disks* on page 9-16 for more information.

To display detailed information about the member, click the other tabs:

- `Enclosure` – Displays details about the array power supplies, cooling fans and temperature. See *Displaying Member Enclosure Status* on page 9-32 for more information.
- `Controllers` – Displays control module and battery status and allows you to set the cache mode. See *Managing Control Modules* on page 9-23 for more information.

- **Disks** — Displays the disk configuration. See *Managing the Disk Array* on page 9-12 for more information.
- **Network** — Displays the network configuration and enables you to configure network interfaces. See *Managing the Member Network Configuration* on page 9-18 for more information.
- **Connections** — Displays the total number of iSCSI connections to the member. See *Displaying Member Connections* on page 9-35 for more information.
- **Service** — Displays detailed array information intended for Technical Support. See *Obtaining Member Service Information* on page 9-36 for more information.

Managing the Disk Array

Each group member includes a highly-available, high-performance disk array. After an array is added to a group, the storage in the member will be available once you set the RAID policy. Most disk array management operations are performed automatically, including RAID configuration and data provisioning.

Disk array management topics include the following:

- *Displaying the Disk Configuration* on page 9-12
- *Changing the Member RAID Policy* on page 9-14
- *Installing and Removing Disks* on page 9-16
- *Handling Failed Disks* on page 9-17


Displaying the Disk Configuration

To display the status of a disk array, click:

```
Members > member_name > Disks tab
```

The Member Disks window appears, displaying the disks in the array. The output shown depends on your array model and the installed disk type (SAS or SATA). Figure 9-9 shows the disks in a PS5000 array.

Figure 9-9: Member Disks - PS5000 Array



The Disk Array Summary panel shows a graphical representation of the disks in the member.

The Installed Disks panel shows detailed information about each disk, including the disk slot, type, model and revision, size, status, and errors.

Note: If disk errors occur, you should investigate before the disk fails.

Disk status is shown in Table 9-3.

Table 9-3: Disk Status

Status	Description
disk-too-small	Disk is smaller than other disks in the array. The disk cannot be used in the array. Note: You should replace the disk with a disk that is the same size as, or larger than, the other disks in the array.

Table 9-3: Disk Status (Continued)

Status	Description
failed	Disk has failed. Note: If the disk status is <code>failed</code> , immediately replace the failed disk. It is recommended that spare disks be installed and functioning in an array.
foreign disk	Disk has a foreign disk label. This indicates that the disk was removed from a different array and then installed in this array. Note: To use the disk in the array, click the <code>foreign disk</code> link and confirm that you want to clear the label.
history-of-failure	Previously failed disk. Note: You may want to consider replacing the disk. If you are sure you want to use the disk in the array, click the <code>history-of-failure</code> link and confirm that you want to use the disk.
offline	Used when a disk does not fall into the other status categories.
online	Functioning as a data disk.
spare	Functioning as a spare disk.
unsupported-version	Disk is not supported by the firmware version on the array.

Select a disk to display the `Disk Statistics` link in the Activities panel. Click this link to display disk statistics.

Changing the Member RAID Policy

Once you select a RAID policy for a member, you can convert the RAID policy only if the new policy requires less disk space than the current policy. For example, you can perform only the following RAID policy conversions:

- RAID-10 to RAID-50
- RAID-10 to RAID-5
- RAID-50 to RAID-5
- RAID-10, no-spares to RAID-50, no-spares (You must use the CLI command `member select member_name raid-policy policy` to convert to a no-spares policy.)

If a member has a RAID policy that cannot be converted, you can change the policy only by removing the member from the group, re-adding it to the group, and then setting the desired RAID policy.

To change the RAID policy for a member, click:

Members > *member_name* > Modify RAID configuration

The Modify RAID Configuration dialog box appears (Figure 9-10), displaying the current RAID policy. As you choose a new RAID policy, the values in the table are updated, showing the new member capacity statistics.

Figure 9-10: Modify RAID Configuration

Modify RAID configuration

Member member2

RAID policy:

☐ RAID-50

☒ RAID-10

☐ RAID-5

After selecting the RAID policy, you can change it only in certain cases as follows:

RAID-10 to RAID-50
RAID-10 to RAID-5
RAID-50 to RAID-5

Expand group capacity:

☒ Immediately expand group capacity

☐ Wait until the member storage initialization completes

RAID-10

	Current	New	Change
Estimated member capacity	3.05 TB	3.05 TB	0 MB
Estimated free space	3.05 TB	3.05 TB	0 MB

OK Cancel

Java Application Window

In addition, if you previously selected to delay using the member's disk space until the RAID verification completes and batteries are fully charged, you can disable this functionality by selecting *Immediately expand group capacity* (recommended setting). If *Immediately expand group capacity* is selected, member space will be immediately available; however, member performance will not be optimal while the RAID verification is in progress.

Click **OK** to complete the RAID policy change. While the RAID policy is changed, the member's RAID status will be *expanding*.

Installing and Removing Disks

You can install disks in an array and replace failed disks without disrupting I/O operations. See the PS Series *Hardware Maintenance* manual for your array model for instructions for physically installing and removing disks.

It is important to note the following:

- **Do not remove a disk that is not a spare.** Removing a disk that is an active member of a RAID set will be handled as a disk failure. If you must remove a disk (for example, to install it in another array), remove a spare from a member that has two working spares if possible. Replace the spare as soon as possible.
- **Do not leave a disk slot empty.** Always install a disk or a blank carrier.
- **Disks are delicate.** Always use electrostatic protection when handling disks. When not installed in an array, always protect a disk from electrostatic discharge.
- **Wait for a disk to spin down.** Once you press the disk handle release button, wait 30 seconds before removing the disk.
- **Use a uniform disk size.** Although you can use disks with different sizes in an array, the smallest disk in the array will determine how much space is available. For example, if an array has 250 GB and 400 GB disks, only 250 GB of space will be available on each disk.

When installing disks in an array, there is a two-minute delay between the time you insert a disk and the time the disk is automatically configured into a RAID set. This time interval enables you to install multiple disks and configure all the disks simultaneously, which is more efficient than installing a single disk, configuring it, and then repeating the process.

For example, when you install a disk in an array, a timer starts. If no other disks are installed, the disk is configured after a delay of 2 minutes. If you install another disk before 2 minutes have elapsed, the timer is restarted.

Notes: If you install a disk during an ongoing RAID reconstruction or verification operation, the new disk will not be configured until the operation completes.

If you remove a disk from one array and then install the disk in a different array, you must clear the disk label before the disk can be automatically configured. Click the `foreign disk` link in the Disk Status window. (See Table 9-3.)

After installing a disk, check the status to ensure that the disk is either online or a spare, as shown in the Member Disks window (Figure 9-9).

Handling Failed Disks

A member's disk space is protected with RAID and one or two spare disks, except in the case of no-spare RAID policies, as described in *Introduction to Member RAID Policies* on page 9-1.

If a member has experienced a disk failure, the Member Disks window (Figure 9-9) shows the failure, and the Alarms panel shows an alarm.

Note: Be sure to replace failed disks as soon as possible. Avoid running with a degraded RAID set. Disk failure in a RAID set that is already degraded will result in loss of data.

When a disk in a RAID set fails, a member behaves as follows:

- **Spare disk is available.** Data from the failed disk is reconstructed on the spare. During the reconstruction, the RAID set that contains the failed disk will be temporarily degraded. After reconstruction, there is no impact on performance.
- **Spare disk is not available, but the RAID set is intact.** The RAID set that contains the failed disk is degraded. For RAID 5, performance is diminished.
- **Spare disk is not available, and the RAID set is degraded.** The member is set offline, and any volumes and snapshots that have data located on the member are set offline. Data may be lost and must be recovered from a backup or replica.

For disk installation information, see *Installing and Removing Disks* on page 9-16.

When you replace a failed disk, a member behaves as follows:

- **Spare disk was used.** The new disk automatically becomes a spare. There is no impact on performance.
- **RAID set was degraded.** Data is automatically reconstructed on the new disk and performance returns to normal after reconstruction.
- **Member was offline because of multiple RAID set disk failures.** Any volumes and snapshots with data on the member will be offline and the data may be lost.

If you have backups, you may be able to recover volume data, as follows:

1. Remove the offline member from the group, as described in *Removing a Member from a Group* on page 9-30. Once removed, the array will be automatically reset to its original condition. This will remove all volume data and group configuration information from the array.

2. Delete the affected volumes. The volume snapshots will also be deleted.
3. Recreate the affected volumes.
4. Recover the volume data from a backup or replica.

After installing a disk, check the status to ensure that the disk is either online or a spare, as shown in the Member Disks window (Figure 9-9).

Managing the Member Network Configuration

A array must have at least one functioning network interface connected to a network and configured with an IP address. Each control module has three Ethernet ports, labeled 0, 1, and 2, for a maximum of three active network connections.

Make sure you have at least one (and preferably three) network connections to the active control module. You should also have at least one (and preferably three) network connections to the secondary control module, in case the active control module fails.

All the normal network rules apply to members in a PS Series group, as described in Table 2-1 on page 2-2.

Note: If you experience network problems, group members may 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 group gets into a state where members lose communication, identify the network problems that caused it and fix them. This will restore the group to normal full operation, including network communication.

Member network configuration topics include the following:

- *Configuring a Network Interface* on page 9-19
- *Displaying Network Interface Status* on page 9-20
- *Enabling or Disabling a Network Interface* on page 9-21
- *Deconfiguring a Network Interface* on page 9-21
- *Modifying the Default Gateway* on page 9-22

Configuring a Network Interface

After connecting a network interface to a network, as described in the *Hardware Maintenance* manual for your PS Series array model, you can configure the interface by assigning it an IP address and subnet mask. You can also modify an existing network interface configuration.


Note: Changing the IP address for a functioning network interface will disconnect all iSCSI initiators that are connected to that interface. However, all initiators should reconnect automatically.

To configure a network interface or modify a network interface configuration, click:


Members > *member_name* > Network tab

The Member Network window appears (Figure 9-11), displaying the current network configuration for each network interface on the member.

Figure 9-11: Member Network



To configure (or modify) a network interface, double-click the network interface in the IP Configuration panel or select the network interface and click **Modify IP settings** in the Activities panel. The Modify IP Settings dialog box appears (Figure 9-12).

Figure 9-12: Modify IP Settings

Specify information about the network interface in the following fields and then click **OK**:

- **IP address** – IP address for the network interface. Be sure to follow the network requirements in Table 2-1 on page 2-2.
- **Subnet mask** – Subnet mask (netmask) for the previous IP address.

Note: The default gateway is the same for all network interfaces on an array. You cannot modify the default gateway using the Modify IP Settings dialog box. See *Modifying the Default Gateway* on page 9-22 for more information.

Once you configure a network interface, you must enable it, as described in *Enabling or Disabling a Network Interface* on page 9-21.

Displaying Network Interface Status

To display detailed information about a member's network interfaces, click:

Members > *member_name* > Network tab

The Member Network window appears (Figure 9-11), displaying detailed information about the network interfaces.

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 is set by administrative action and can be enabled or disabled.
Note: If the operational status is down, but the requested status is enabled, you should identify and correct the error.
- Speed. Ensure that the interface speed is as expected.
- MTU size. Each interface automatically determines the path MTU size. You can set the Ethernet frame size on an iSCSI initiator.
- Packet errors. A few packet errors may not indicate a problem. A large number of packet errors should be investigated and may indicate a network problem.

Note: If a port or network interface fails or if a large number of packet errors occur, you may be able to regain network connectivity by moving the cable connection to a functioning network interface port on the active control module.

The IP Configuration panel shows each interface and its assigned IP address, netmask, MAC address and description, if any.

Enabling or Disabling a Network Interface

Disabling a member's network interface makes it unable to service network I/O requests. Enabling an interface makes it operational if it is configured properly.

Note: Disabling a functional network interface will disconnect all iSCSI initiators that are connected to that interface.

To enable or disable a network interface, click:

Members > *member_name* > Network tab

The Member Network window appears (Figure 9-11).

In the IP Configuration panel, select the network interface. In the Activities panel, click `Enable interface` to enable the interface or click `Disable interface` to disable the interface.

Deconfiguring a Network Interface

You can deconfigure a network interface so that it is no longer used for I/O. Alternately, you can temporarily disable the interface, as described in *Enabling or Disabling a Network Interface* on page 9-21.

To deconfigure a network interface, click:

Members > *member_name* > Network tab

The Member Network window appears (Figure 9-11).

In the IP Configuration panel, select the network interface and then click **Modify IP Settings** in the Activities panel. The Modify IP Settings dialog box appears (Figure 9-12).

In the IP address field, delete the IP address. Then, click **OK**.

Modifying the Default Gateway

When you add a member to a group, the `setup` utility prompts you for the default gateway that will be used for all the member's network interfaces.

You must specify a default gateway for a member if you want the IP addresses assigned to the member's network interfaces to be able to communicate outside of the local network (for example, for host access to volumes).


To modify the default gateway for a member, click:

Members > *member_name* > Network tab

The Member Network window appears (Figure 9-11).

In the IP Configuration panel, click **Modify** next to the `Default gateway` field. The Set Default IP Gateway dialog box appears (Figure 9-13).

Figure 9-13: Set Default IP Gateway



Specify the default gateway in the dialog box or delete the IP address if a default gateway is not desired. Then, click **OK**.

Managing Control Modules

Each array has one or two control modules installed. Only one functioning control module is required for operation.

Caution: Only install control modules of the same type in an array. Do not mix control module types in an array.

When a dual control module array starts, one control module is designated as active (responsible for serving I/O to the member), while 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. The active control module is designated by the green LED labeled ACT.

Under normal operation, the status of a control module (active or secondary) will not change, unless the member is restarted. However, if the active control module encounters a failure, the secondary control module will take over and serve I/O. This is called control module failover.

Control module management topics include the following:

- *Displaying Control Module Status* on page 9-23
- *Setting Control Module Cache Modes* on page 9-25
- *Handling Failed Control Modules* on page 9-27


Displaying Control Module Status

To display the control module status, click:

Members > *member_name* > Controllers tab

The Member Controllers window appears (Figure 9-14).

Figure 9-14: Member Controllers



Each Control Module Slot panel shows the following information:

- Status
- Boot time
- Cache and NVRAM battery status, if applicable to the control module model
- Model number
- Boot ROM version
- PS Series firmware version

An empty slot means that a control module is not installed or has failed.

Note: If a control module fails, see *Handling Failed Control Modules* on page 9-27 for more information.

Control module cache battery status is shown in Table 9-4. If the cache battery charge drops below acceptable limits and the low-battery-safe policy is enabled, the active cache

will automatically switch to write-through mode. Note that write-through mode may have a negative impact on performance.

Table 9-4: Cache Battery Status

Status	Description
good	Battery is fully charged.
failed	Battery has failed. Note: If the battery status is failed or missing, or if the status is low voltage or low voltage, charging status for an extended period of time, contact your array support provider for information on how to proceed.
missing	Battery is missing.
low voltage	Battery is below the limit for normal operation.
low voltage, charging	Battery is charging but is still below the limit for normal operation.
good, charging	Battery is charging but has enough charge for normal operation.

NVRAM coin cell battery status can be one of the following:

- good – Battery is fully charged
- failed – Battery has failed

If the NVRAM battery on the active control module fails, the active control module will fail over to the secondary. If the array has only one control module and the battery fails, the write cache will operate in write-through mode.

Note: For information on replacing a battery, contact your array support provider.

Setting Control Module Cache Modes

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

- **Write-back.** Data is stored in the cache until it is written to disk, which may improve performance. The cache battery protects the data. In a dual control module configuration, when write-back mode is enabled, cache data is also mirrored across the two controllers, providing additional data protection.
- **Write-through.** Data is written immediately to disk. This reduces the risk of data loss, but may have a negative impact on performance.

You can set policies that control whether a cache will operate in write-back or write-through mode when a cache battery voltage is low or when a control module fails in a dual control module configuration.

Note: It is recommended that you use the default cache policies, described below.

To display the control module cache policies, click:

Members > *member_name* > Controllers tab

The Member Controllers window appears (Figure 9-14).

In the Memory Cache panel, choose the desired cache mode policies:

- **Single-controller-safe policy.** To enable, select `Use write-through mode if only one controller is functional`.

If enabled on a member with a single control module, the control module cache will always operate in write-through mode. If disabled (the default), under normal conditions, the cache will operate in write-back mode.

If enabled on a member with dual control modules, but only one module is functioning, the cache on the functional control module will operate in write-through mode. If disabled (the default), the cache on the functional control module will continue to operate in write-back mode.

- **Low-battery-safe policy.** To enable, select `Use write-through mode if battery charge is below tolerance`.

If enabled (the default) on a member with a single control module, and the charge on the cache battery is low, the cache will operate in write-through mode.

On a member with dual control modules, if the low-battery-safe policy is enabled (the default) and the single-controller-safe policy is disabled, the active control module cache will continue to operate in write-back mode if the cache battery has an adequate charge, even if the secondary control module's cache battery has a low battery charge. If both the low-battery-safe and single-controller-safe policies are enabled, the active control module cache will operate in write-through mode if the cache battery charge on either control module is low.

If the low-battery-safe policy is disabled, a cache will operate in write-back mode, regardless of cache battery charge. Disabling low-battery-safe mode is not recommended because of the risk of data loss.

After choosing the cache policies, click `Save changes`. If you make a mistake, click `Discard changes`.

Handling Failed Control Modules

The member behavior when a control module is removed or fails depends on the number and type of control modules, the network cabling configuration, and the cache mode settings. See the *Hardware Maintenance* manual for your array model for detailed information about array-specific control module and network failover behavior.

In a dual control module array, if the secondary control module is removed or fails, the remaining control module may enter write-through mode depending on the cache mode policies.

In a dual control module array, if the active control module fails, the secondary control module automatically takes over and become active. If there is a network connection to the active control module, the control module failover will be transparent to applications; however, iSCSI initiators must reconnect to the group IP address.

If the only functioning control module fails, the member will be inaccessible from the network and you may experience data loss.

If a control module fails, replace it as soon as possible with the *same* control module type.

Caution: Do not mix control module types in an array.

For information about replacing a control module, see the *Hardware Maintenance* manual for your array model or contact your array support provider. Do not remove a failed control module until you have a replacement.

Moving a Member Across Pools

In a multi-member group, you can separate storage space into multiple storage pools and assign members to the pools. A member can be in only one pool at a time. Volumes can be assigned to pools with sufficient free space; volume data will be restricted to the pool members.

When a member is first added to a group, by default, it is assigned to the default storage pool. If multiple pools are available, you can assign the member to a different pool.

You can also move a member from one pool to another, with no impact on users and data availability. Volume data located on the member will be moved to the remaining pool members.

Within a pool, it is recommended that 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, both members should be configured with the same RAID level.

Notes: When moving a member to a different pool, if the remaining pool members do not have sufficient space to store the volume data from the member, the move operation will not start.

The space in a moved member will not be available in the new pool until the move operation completes.


If you cancel an in-progress member move operation, the member space will immediately return to the original pool.

If the pool to which a member is assigned is deleted, the members will be automatically and instantly assigned to the default pool.

To move a member to a different storage pool, click:

Members > *member_name* > Modify member settings

The Modify Member Settings dialog box appears (Figure 9-15), displaying the available pools and the capacity and free space in each pool.

Figure 9-15: Modify Member Settings


Modify member settings

Member member2

General properties:

* Member name:

Description:

Storage pool assignment:

Storage Pool	Capacity	Free	Drives
☐ default	7.85 TB	6.52 TB	SATA
☒ user	3.05 TB	3.05 TB	SATA

[Create new pool](#)

Storage pool user

	Current	New	Change
Total pool space	3.05 TB	3.05 TB	0 MB
Free pool space	3.05 TB	3.05 TB	0 MB

OK Cancel

Java Application Window

Select the pool for the member. The values in the storage pool table will be updated. If there is sufficient space for the member move operation, click **OK**.

The time to complete a member move depends on the amount of volume data that must be moved to the remaining pool members.

You can cancel a member move operation, if desired, and return volume data to the original storage pool. To cancel a volume move operation, click:

Members > *member_name* > Cancel member move

Modifying a Member Name and Description

You can modify a member name and description. The name is used only to identify the member for administrative purposes or replication.

To rename a member, click:

Members > *member_name* > Modify member settings

The Modify Member Settings dialog box appears (Figure 9-15).

Enter the new member name. The name must be unique in the group. See Table 2-2 for information on member name requirements and restrictions.

Optionally, specify a description for the member.

Click **OK** to apply the changes.

Modifying the Group Membership Password

When you create a group with the `setup` utility, you are prompted for the password that will be required when you add a member to the group.

To modify the group membership password, click:

Members > Set password

The Change Group Membership Password dialog box appears (Figure 9-16). Enter the new password and confirm the password. The password can contain 3 to 16 alphanumeric characters and is case-sensitive. Then, click **OK**.

Figure 9-16: Change Group Membership Password



Removing a Member from a Group

You can remove a member from a multi-member group while the group is online and data remains available. You cannot remove the only member from a group.

After a member is removed from the group, it is automatically reset, which eliminates all group, member, and volume information and any volume data residing on the member.

Removing a member will decrease the overall storage capacity of the group and also the storage pool to which the member is assigned. Any volume data located on the removed member must be vacated and moved to the remaining pool members.

If you attempt to remove a member and the remaining pool members do not have enough space to hold the volume data, you will receive an error message telling you that there is not enough space, and the remove operation will not start until you free enough pool space.


The time to complete a member removal depends on the amount of volume data that must be moved to the remaining pool members. While data is being moved, the member status is *vacating*.

To remove an online member from the group, click:

Members > *member_name* > Delete member

The Delete Member dialog box appears (Figure 9-17). Click *yes* to confirm that you want to delete the member.

Figure 9-17: Delete Member



In some cases (for example, if the member has experienced a catastrophic failure), you must remove a member that is offline. Follow the same procedure as for removing an online member. However, an error message will appear indicating that the vacate has failed, which means that the data could not be moved to the other pool members. When prompted, confirm that you want to remove the member.

Using LEDs to Identify a Member

In a lab environment, it may be necessary to identify a member for maintenance or service purposes. To help you locate a member, you can make the LEDs on a member chassis flash.

For example, on the front panel of a PS400E array, the ALRM and WARN LEDs will flash. In addition, on the back panel of any array model, the fan tray LED and the control module ERR LED will flash.

To flash a member's LEDs, click:

Members > *member_name* > Start LEDs flashing

To stop flashing a member's LEDs, click:

Members > *member_name* > Stop LEDs flashing

Displaying Member Enclosure Status

To display information about a member's power supplies, cooling fans, and temperature, click:

Members > *member_name* > Enclosure tab

The Member Enclosure window appears. The output depends on your array model. Figure 9-18 shows output from a PS5000 array.

Figure 9-18: Member Enclosure - PS5000 Array

The screenshot shows the 'group1 Group Manager' application window. The top bar indicates the user is logged in as 'gradmin' at '4/10/07 1:47:13 PM'. The main window is titled 'Member member1' and has several tabs: Status, Enclosure (selected), Controllers, Disks, Network, Connections, and Service. On the left is a sidebar with a tree view containing 'Group group1', 'Group Configuration', 'Group Monitoring', 'Events (2 new)', 'Storage Pools', 'Members' (selected), 'Volumes', 'Volume Collections', and 'Replication Partners'. Under 'Members', 'member1' is highlighted. The main content area under the 'Enclosure' tab is divided into three sections: 'Power Supplies', 'Cooling Fans', and 'Temperature Sensors'. Each section contains a table of data.

Power Supplies

Power supply	Status
Power supply 0	OK
Power supply 1	OK

Cooling Fans

Cooling fan	Status	Speed, RPM	Normal RPM range
Power Cooling Module 0 Fa...	normal	4920	2100 - 5200
Power Cooling Module 0 Fa...	normal	4800	2100 - 5200
Power Cooling Module 1 Fa...	normal	4770	2100 - 5200
Power Cooling Module 1 Fa...	normal	4770	2100 - 5200

Temperature Sensors

Sensor	Status	Temperature	Normal range
Control module 0 processor	normal	64° C	10° - 90° C
Control module 0 chipset	normal	32° C	10° - 45° C
Control module 1 processor	normal	60° C	10° - 90° C
Control module 1 chipset	normal	32° C	10° - 45° C
Control module 0 SAS Controller	normal	60° C	10° - 110° C
Control module 0 SAS Expander	normal	66° C	10° - 110° C
Control module 0 SES Enclosure	normal	33° C	10° - 45° C
Control module 1 SAS Controller	normal	61° C	10° - 110° C

At the bottom of the window, there is a status bar showing '1 outstanding alarm (1 new)'.

The Power Supplies panel shows the status of the array power supplies. Power supply status is shown in Table 9-5.

Table 9-5: Power Supply Status

Status	Description
OK	Power supply is providing power to the array.
no AC power or missing	Power supply is not installed, not turned on, or not connected to a power source. Be sure to keep all power supplies installed and turned on.
failed	Power supply has failed.

Some power supplies have a fan (only in PS Series Model 50E to 400E arrays). Power supply fan status is shown in Table 9-6.

Table 9-6: Power Supply Fan Status

Status	Description
OK	Fan is functioning.
fan failed	Fan has failed, or the power supply is not installed, not turned on, or not connected to a power source.

Note: If a power supply or power supply fan fails, replace the power supply as soon as possible. To ensure 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 array support provider.

One fan failure will not affect the member. However, more than one fan failure will cause the member temperature to rise. If the temperature rises excessively, an event message is generated to warn administrators of the potential problem. If the member temperature reaches its maximum limit, event and alarm messages are generated.

The Cooling Fans panel provides information on the fan tray or cooling module fans. Fan status is shown in Table 9-7.

Table 9-7: Fan Status

Status	Description
normal	Fan is operational.
warning	Fan RPMs are outside the normal range of operation.
critical	Fan has failed.

Note: If a fan tray or cooling module fan fails, replace the affected component before the array overheats and is damaged. In addition, if fan RPMs are outside the normal range, the fan may be at risk of failure, and you may want to replace the fan tray or cooling module. For proper cooling, do not remove a fan tray or cooling module until you have a replacement. For information about replacing hardware see the *Hardware Maintenance* manual for your array model or contact your array support provider.

The Temperature Sensors panel displays the temperature for the various array controllers and processors, in addition to the normal temperature range.

Array temperature status is shown in Table 9-8.

Table 9-8: Array Temperature Status

Status	Description
normal	Temperature is within normal range.
warning	Temperature is outside normal range, but within limits.
critical	Temperature is outside operating limits.

Note: If an array temperature status is `critical`, identify and correct the problem before the array overheats and is damaged. For example, check the air conditioning system where the array is located, in addition to the air flow around the array. If a processor temperature remains high, you may need to replace the control module. For proper cooling, do not remove a control module until you have a replacement. For information about replacing a control module, see the *Hardware Maintenance* manual for your array model or contact your array support provider.

PS Series Model 50E to 400E arrays also display the system (ambient) temperature, which is obtained using the two sensor temperatures (in Celsius) from the control module with the highest temperatures and the following calculation:

$$((\text{Backplane Sensor 0} + \text{Backplane Sensor 1}) / 2) - 7$$


Displaying Member Connections

To display information about the connections to a member, click:

Members > *member_name* > Connections tab

The Member Connections window appears (Figure 9-19).

Figure 9-19: Member Connections



The iSCSI Connections panel displays information about the initiator address, which volume it is connected to (Target column), how long the connection has been active, and which Ethernet port the initiator is using.

Obtaining Member Service Information

At some point, you may need to provide detailed information about member hardware (for example, a component model, revision, or serial number) to your array support provider. To display information about the hardware, click:

Members > *member_name* > Service tab

The Member Service window appears. The output shown depends on your array model.

Figure 9-20 shows the output from a PS5000 array.

Figure 9-20: Member Service - PS5000 Array

The screenshot shows the 'group1 Group Manager' application window. The 'Member member1' service window is active, displaying hardware information for a PS5000 array. The interface includes a left sidebar with navigation options like Group Configuration, Group Monitoring, Events, Storage Pools, Members, Volumes, Volume Collections, and Replication Partners. The main area shows the 'Service' tab with two tables: 'Component Versions' and 'Disk Versions'.

Component	Part number	Revision	Serial number	ECO level
Backplane	62532	NA	SHU75212000038D	NA
PCMD	64362	NA	PMC643620137337	NA
PCM1	64362	NA	PMC643620137336	NA
CM0	70-0111	R05	RMS8330800010ED	C00
CM1	70-0111	R05	RMS8330800010DB	C00

Slot	Type	Model	Revision	Serial number
0	SAS	MODEL1234	3.123	9QG0GQGG
1	SAS	MODEL1234	3.123	9QG0GJRD
2	SAS	MODEL1234	3.123	9QG0GQLV
3	SAS	MODEL1234	3.123	9QG0F2L5
4	SAS	MODEL1234	3.123	9QG0GJMN
5	SAS	MODEL1234	3.123	9QG0GJ9G
6	SAS	MODEL1234	3.123	9QG0GJ9L
7	SAS	MODEL1234	3.123	9QG0GJX0

The bottom status bar indicates 'No outstanding alarms'.

10 Diagnosing Problems

Topics on how to interpret status and diagnose problems in a group include:

- *Monitoring Events* on page 10-1
- *Monitoring the Group* on page 10-7
- *Monitoring Performance Statistics* on page 10-15
- *Using MRTG to Monitor I/O and Performance* on page 10-18

Caution: If a problem occurs in a PS Series array, it is strongly recommended that you do *not* turn off power to the array unless the array has been cleanly shut down. Use the `shutdown` command to cleanly shut down an array.

Monitoring Events

When an event occurs in a group (such as the creation of a volume or the removal of a power supply), a message is generated. These messages enable you to track normal operations and detect problems before they disrupt operations. Event messages can also help you diagnose problems.

Monitoring events is described in the following topics:

- *Displaying Events* on page 10-1
- *Accessing the Event Log File on a Remote System* on page 10-4
- *Accessing Events Sent to an E-Mail Address* on page 10-4
- *Examining SNMP Traps* on page 10-4
- *Monitoring Alarms and Action Items* on page 10-4

Displaying Events

Events can be displayed using the Group Manager GUI or CLI and are also seen on the console. Each event is assigned a priority. Table 10-1 lists event priorities in order from highest (most severe) to lowest (least severe).

Table 10-1: Event Priorities

Priority	Description
FATAL	Catastrophic failure that requires immediate administrator intervention.
ERROR	Serious failure that should be investigated and may require administrator intervention.
WARNING	Potential problem in the group. An event with this priority may become an event with <code>error</code> priority if administrator intervention does not occur.
INFO	Normal operational and transitional events that are of interest to administrators but do not require intervention.

Examples of event messages are shown next:

```
INFO      4/11/06 11:02:23 mem2 Volume 'dbvol' has been successfully created
```

```
WARNING 4/11/06 14:19:43 mem1 There are outstanding warning health conditions
```


Each event message displays:

- Event priority
- Date and time that the event occurred
- Member that generated the event
- Event text

You can view events by clicking `Events` in the far left panel of the GUI. The Group Events window appears (Figure 10-1).

To view additional information about events, place the cursor over the event message. Detailed information appears in a popup window. If the popup does not appear, check that you have enabled popups, as described in *General Policies* on page 2-12.

Figure 10-1: Group Events



From the Group Events window, you can:

- **View all events or events of a specific priority.** Pull down the **View** menu and select the events you want to display, either all events, warning and error events, or fatal (critical) events.
- **Retrieve previous events.** To retrieve the most recent 100 events, click the left-pointing arrow (**More**) icon. Click the icon 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 double checkmark () icon. Acknowledging all events turns off the bolding.
- **Clear the event list.** To erase all the events from the panel, click the X () icon. To display the events again, click the retrieve previous events icon.
- **View and hide details.** Click the magnifying glass icons () to display and hide detailed information about the selected event.

If the GUI is not available, you can use the Group Manager CLI to display events. Use telnet or SSH to connect to the group and log in to a group administration account. If you do not have a network connection, use a serial connection from an array to a console or computer running a terminal emulator. At the CLI prompt, enter the `show recentevents` command. See the *CLI Reference* for more information.

Accessing the Event Log File on a Remote System

If you enabled syslog notification, as described in *Setting Up Event Notification* on page 3-20, you can examine the file to which events are logged on the remote host. To do this, log in to one of the remote hosts. The location of the log file depends on how syslog is configured on the remote host.

Accessing Events Sent to an E-Mail Address

If you enabled e-mail notification of events, as described in *Setting Up Event Notification* on page 3-20, you can access the event messages that are sent to the configured e-mail addresses.

When an event is generated, a timer starts. If no other events are received within one minute, e-mail is sent to the addresses you configured for notification. If another event is received within one minute, the timer is restarted and the process is repeated. If the timer continues to restart, after 2 minutes, e-mail is sent.

Examining SNMP Traps

If you enabled SNMP trap notification, as described in *Using SNMP to Monitor a Group* on page 3-31, you can examine any traps sent from the group. You will need an SNMP console to examine traps.


Monitoring Alarms and Action Items

The Group Manager GUI notifies you when an alarm occurs in the group. An alarm is an event that indicates a persistent condition in an array (for example, high temperature). Alarms enable you to detect problems before they disrupt operations. Each alarm is assigned a priority level, either warning or critical, depending on the severity of the problem.

The GUI also indicates when there is an in-progress action item that requires an administrator to complete (for example, manual replication transfer).

When a new alarm occurs, the Outstanding Alarms panel header, located at the bottom of each main GUI window, will flash until you click the panel and open it (Figure 10-2).

Figure 10-2: Outstanding Alarms



The Outstanding Alarms panel includes three tabs that show:

- Critical alarms, as described in Table 10-2
- Warning alarms, as described in Table 10-3
- Action Items to be completed

Each alarm entry displays the priority, the member that reported the alarm, and the message text. Place the mouse over the message text to display detailed information. You can sort the information by clicking the table headers.

Note: Alarms that correspond to member health conditions also light LEDs on the array.

Table 10-2: Critical Alarms

Priority	Description
CRITICAL	One or more of the following has occurred: • RAID is not functioning. • Lost block table is full. • Array temperature exceeds upper or lower limit. • Control module cache has lost data. • One or both fan trays are not installed. • Both fans on a fan tray or cooling module have failed. • Cache battery has less than 72 hours of charge or is not being charged because it exceeds the temperature limit. • NVRAM coin cell battery has failed. • Cache contains data that does not belong to any of the installed disks. • More than one valid RAID set exists in the array. • Control modules are different models. • A critical hardware component has failed. • Operations panel failed or not installed (only PS5000 arrays). • Processor that monitors array has failed (only PS5000 arrays).

Table 10-3: Warning Alarms

Priority	Description
WARNING	One or more of the following has occurred: • RAID set is degraded but still functioning. • RAID (volume-level) has lost blocks. • Component temperature is near upper or lower limit. • Fan RPMs exceed upper or lower limit. • Power supply fan has failed. • Power supply is not installed or has no power. • Only one control module installed or control module has failed over. • Control module has insufficient amount of RAM. • Lock on secondary control module is open (only PS50E to PS2400E arrays). • Active control module syncing with secondary. • No communication between control modules. • Installed spare does not have enough capacity to replace a RAID set disk. • A non-critical hardware component has failed. • Real-time-clock battery is low.

Alarms are displayed until the condition is corrected. However, the event message associated with the alarm will remain in the event log. Click the double checkmark icon in the Outstanding Alarms panel to acknowledge all alarm and turn off bolding of any outstanding alarm messages.

Action items are displayed until you complete the task indicated.

Close the Outstanding Alarms panel by clicking the header.

Monitoring the Group

Normally, a PS Series group requires little monitoring. Regularly examining the events will ensure that you deal with issues promptly before they result in problems.

Group monitoring topics include the following:

- *Using the Group Monitoring Windows* on page 10-7
- *Monitoring Space* on page 10-12
- *Monitoring Replication* on page 10-13
- *Monitoring Volumes and Snapshots* on page 10-14
- *Monitoring Hardware* on page 10-14

Using the Group Monitoring Windows


The GUI includes windows that enable you to monitor current group activity, including iSCSI connections to the group iSCSI targets, snapshot and replica schedules, replication activity, administrative login activity and history, and in-progress operations.

Displaying iSCSI Connections

To show current iSCSI connections in the group, in the far left panel, click:

Group Monitoring > iSCSI Connections tab

The Group Monitoring - iSCSI Connections window appears (Figure 10-3).

Figure 10-3: Group Monitoring - iSCSI Connections


The iSCSI Connections tab displays the iSCSI connections to iSCSI targets (volumes and snapshots), including the initiator address and the megabytes read and written since the last iSCSI connection.

Displaying Schedules

To show current snapshot and replica schedules, in the far left panel, click:

Group Monitoring > Schedules tab

The Group Monitoring - Schedules window appears (Figure 10-4).

Figure 10-4: Group Monitoring - Schedules

The Schedules tab displays the volumes and their schedule configurations. In the Actions column, click the icons to modify, enable or disable, or delete a schedule.

Displaying Replication Information

To monitor replication in the group, in the far left panel, click:

Group Monitoring > Replication tab

The Group Monitoring - Replication window appears (Figure 10-5).

Figure 10-5: Group Monitoring - Replication

group1 Group Manager Account: grpadmin Logged in 3/19/07 11:47:43 AM Logout

Group Monitoring

Summary

ICS Connections
ICS connections: 4

Snapshot Schedules
Volume schedules: 2
Collection schedules: 0

Replication Schedules
Volume schedules: 1
Collection schedules: 0

Volume Replication
Outbound replication: 7
In-progress replication: 1
In-progress replication: 4

Administrative Sessions
Active GUI sessions: 1
Active CLI sessions: 0

In-Progress Operations
Member operations: 0
Volume operations: 0

Configured Volume Replication

Direction	Volume	Status	Started	Transferred
group1 ↔ group2	vol-1	ready	2/6/06 9:28:16 PM	
group1 ↔ group2	MRvol	manual-in-progress	3/14/07 1:48:38 PM	150 MB / 15...
group1 ↔ group2	testvol	ready		
group1 ↔ group2	vol1	ready		
group1 ↔ group2	vol2	ready		
group1 ↔ group2	vol3	manual-in-progress	3/16/07 12:55:31 PM	150 MB / 15...
group1 ↔ group2	vol3	ready		

Outbound Replication History

Started	Volume	Partner	Duration	Data size	Speed	Transferred
3/16/07 6:36:2...	testvol	group2	2 hr 7 min 7...	705 MB	5 MB/min	complete
3/7/07 9:39:50...	vol3	group2	15 min 46 sec	75 MB	4 MB/min	complete
3/7/07 7:33:22...	vol1	group2	31 min 17 sec	300 MB	9 MB/min	complete

Tools 1 outstanding alarm (1 new)

The Configured Volume Replication panel displays the volumes currently configured for replication, the direction of replication, the status, the start time, and the amount of data transferred.

For each volume, only the last 10 replication operations are shown in the Outbound Replication History panel.

Displaying Login Information

To show administration session information, in the far left panel, click:

Group Monitoring > Administrative Sessions tab

The Group Monitoring - Administrative Sessions window appears (Figure 10-6).

Figure 10-6: Group Monitoring - Administrative Sessions

The screenshot shows the 'Group Manager' application window. The 'Group Monitoring' section is active, and the 'Administrative Sessions' tab is selected. The 'Summary' panel on the left shows various system metrics. The main area displays two tables: 'Active Administrative Sessions' and 'Administrative Login History'.

Active Administrative Sessions

Account	Session type	Login time	Client IP	Local IP
grpadmin	Secure shell (SSH)	4/17/07 1:32:38 PM		
grpadmin	Secure GUI session	4/17/07 2:00:34 PM	172.17.30.22	172.19.106.140
grpadmin	Secure GUI session	4/17/07 2:19:40 PM	172.16.1.112	172.19.106.140
staff23	Secure GUI session	4/17/07 2:18:54 PM	172.16.1.199	172.19.106.140

Administrative Login History

Account	Authentication	Account type	Managed pools	Status	Last Login
staff23	remote	pool administrator (group-read-access)	userpool	enabled	4/17/07 2:18:54 PM
newadmin	local	group administrator	all pools	enabled	
grpadmin	local	group administrator	all pools	enabled	4/17/07 2:19:40 PM

The Active Administrative Sessions panel displays the active GUI and CLI sessions and the login time. The Client IP column shows the IP address of the system running the GUI. The Local IP column shows the IP address to which the system is connected (usually the group IP address).

The Administrative History panel displays the login history for the accounts that have logged in to the group, the account type and authentication, the pools for which they have management privilege, status, and last login time.

Displaying Member and Volume Moving Operations

You can display in-progress operations for members (removing a member from the group or moving a member to a different pool) and volumes (moving a volume to a different pool).


To show in-progress member and volume operations, in the far left panel, click:

Group Monitoring > Operations tab

The Group Monitoring - Operations window appears (Figure 10-7).

If desired, you can cancel the in-progress operation by clicking [Cancel](#).

Figure 10-7: Group Monitoring - Operations



Monitoring Space

You can monitor usage of the following:

- **Group space.** If space is low, consider increasing the group capacity. See *Displaying the Group Configuration and Status* on page 3-1 for information about monitoring group space.

To increase the amount of free space in the group, see *Adding Members to an Existing Group* on page 2-5.

- **Pool space.** If space is low, consider increasing the pool capacity. See *Displaying Storage Pool Configuration and Status* on page 4-6 for information about monitoring pool space.

To increase free pool space, you can move volumes from the low-space pool to a different pool. You can also add more members to the pool, as described in *Moving a Member Across Pools* on page 9-27.

- **Snapshot reserve.** If space is low, consider increasing the snapshot reserve. See *Displaying Snapshot Configuration and Status* on page 6-6 for information about monitoring the snapshot reserve.

To increase the snapshot reserve for a volume, see *Modifying Snapshot Settings* on page 6-3.

- **Delegated space.** If space is low, consider increasing the delegated space. See *Displaying Replication Activity Between Partners* on page 7-32 for information about monitoring delegated space.

To increase delegated space for a partner, see *Modifying a Replication Partner* on page 7-17.

- **Replica reserve.** If space is low, consider increasing the replica reserve. See *Displaying Inbound Replication Status* on page 7-38 for information about monitoring the replica reserve.

To increase the replica reserve for a volume, see *Configuring a Volume or Volume Collection for Replication* on page 7-19.

Monitoring Replication

If there is no progress in a replication operation, a warning alarm will occur. If this occurs or if replicating a volume takes a long time, be sure you have adequate network bandwidth between the groups, in addition to full IP routing. A slow network link between the partners may result in long replication times.

Examine the status of the volume replication and check to see how much data is being moved. Ensure that the network link can handle this much data. You may need to increase the network bandwidth.

See *Displaying the Replication Configuration for a Volume* on page 7-29 and *Displaying Inbound Replication Status* on page 7-38 for information about monitoring replication duration.

Monitoring Volumes and Snapshots

A volume can be set offline if network problems occur or if a member that contains volume data is shut down. If a volume is set offline because of a problem, all its snapshots will be set offline.

See *Displaying Volume Status and Configuration* on page 5-11 for information about monitoring volumes. See *Displaying Snapshot Configuration and Status* on page 6-6 for information about monitoring snapshots.

Monitoring Hardware

Although array hardware is redundant, be sure to monitor members and quickly replace any failed hardware. Hardware problems will result in alarms that also trigger array LEDs.

For information about monitoring the overall member status, see *Displaying Member Configuration and Status* on page 9-6.

Component failures appear in the Member Status windows (Figure 9-7 and Figure 9-8).

You can further monitor array components:

- **Disks and RAID status.** It is important to detect failed disks and replace them as soon as possible. Although spare disks and RAID protect data against disk failures, multiple disk failures may put data in jeopardy. If disk errors occur, consider replacing the disk before it fails completely.

See and *Displaying the Disk Configuration* on page 9-12 for more information.

- **Control modules and batteries.** If the battery status remains at low voltage or low voltage, charging status for an extended period of time, the battery may need to be replaced.

See *Displaying Control Module Status* on page 9-23 for information about monitoring control modules. For information on replacing a cache battery, contact your array service provider.

- **Network interfaces.** If a network interface shows errors, examine the group event log to identify the problem. If a port or network interface fails, you may be able to regain network connectivity by moving the cable connection to a functioning network interface port.

See *Displaying Network Interface Status* on page 9-20 for information about monitoring network interfaces.

- **Power.** Make sure there are no red LEDs on the power supplies and no fan problems (if applicable to the power supply).

See *Displaying Member Enclosure Status* on page 9-32 for information about monitoring power supplies.

- **Cooling.** Periodically, check the room temperature where the hardware is located and ensure that the room is sufficiently cool and ventilated. Make sure there are no red LEDs on the array fan trays or cooling modules and also check member temperature.

One fan failure will not affect a member. However, more than one fan failure will cause the member temperature to rise. If the temperature rises excessively, an event message is generated to warn administrators of the potential problem. If the member temperature reaches its maximum limit, event and alarm messages are generated. Replace the affected fan tray or cooling module as soon as possible.

See *Displaying Member Enclosure Status* on page 9-32 for information about monitoring cooling modules.

Monitoring Performance Statistics

You can use the Performance Monitor to display performance statistics for a member. The Performance Monitor collects statistical data every second.

To start the Performance Monitor, click:

`Members > member_name`

The Member Status window appears (Figure 9-7). In the Activities panel, click the desired component, either disks, network interfaces, IP counters, TCP counters, or UDP counters. The Performance Monitor window appears.

You can also expand the Tools panel in the lower-left corner of any main GUI window and click `Performance monitor` to display the Performance Monitor.


In the far left panel of the Performance Monitor window, select whether to show rates or show values.

On the left side of the task bar above the statistics panel, click the first arrow to start the data polling. Click the double-bar icon to stop polling. Once you stop polling, you can click the other arrows to navigate through the data.


On the right side of the task bar above the statistics panel, click the icons to display the statistics in a chart view, histogram view (bar graph), linear scale and logarithmic scale and to scale the display.

For example, if you click Disks in the Activities panel of the Member Status window and then choose to display a histogram view, the Performance Monitor - Disk Statistics window appears (Figure 10-8), showing disk statistics, including total megabytes transferred to the member and megabytes for each disk.

Figure 10-8: Performance Monitor - Disk Statistics Histogram



To change the displayed statistics, click `Select Statistics`. The Select Statistics dialog box appears (Figure 10-9).

Figure 10-9: Select Statistics

Expand the member name and double-click a component to display more statistics. To display detailed information about a disk, expand the disk and double-click `Disk statistics`. The Performance Monitor - Disk Statistics window appears (Figure 10-10), displaying I/O transfers, megabytes transferred, and errors.



Figure 10-10: Performance Monitor - Disk Statistics

Figure 10-11 shows an example of member network statistics displayed in a chart view, including the bytes transferred and errors for each network interface.

Figure 10-11: Performance Monitor - Network Statistics - Chart View



Using MRTG to Monitor I/O and Performance

As a supplement to the Performance Monitor feature of the PS Series array, you can also use Multiple Router Traffic Grapher (MRTG) to monitor the I/O activity, latency, and throughput of volumes and members. For more information on using MRTG, see the *CLI Reference*.

A Glossary

access control record

Controls host access to a volume and its snapshots. Access can be restricted through IP address, iSCSI initiator name, or CHAP user name.

active control module

In a dual control module array, the control module that is actively serving I/O. To physically identify the active control module, examine the array front panel. The LED labeled ACT will be green on the active control module.

array

A PS Series array is a completely self-contained storage unit that includes at least seven and up to fourteen disks in a highly-available disk array with single or dual control modules and redundant power supplies and fan modules.

bandwidth

Rate at which an I/O subsystem or component can transfer bytes of data. Also, sometimes called the *transfer rate*.

CHAP

Acronym for Challenge Handshake Authentication Protocol, used in conjunction with access control records to restrict host access to volumes and snapshots. CHAP is a network login protocol that uses a challenge-response mechanism. With CHAP authentication, volume and snapshot access can be restricted to hosts that supply the correct account name and password combination.

CLI

Acronym for Command Line Interface, used to manage groups and arrays. CLI commands can be used in scripts.

cloning

Functionality that enables the creation of an exact copy of a volume or a copy of a volume as it existed at the time a snapshot or replica was created.

custom snapshot collection

One or more snapshots created at the same time through a multi-volume snapshot operation.

default gateway

Local router that is used to forward network traffic beyond the local network.

demote

Converts a replicated volume to a failback replica set for the purpose of failing back data from the secondary group, or converts a recovery volume to an inbound replica set, returning it to its original state. Also, see *promote* on page A-6 and *failback* on page A-2.

discovery

Discovery, or auto-discovery, involves an iSCSI initiator requesting from a target portal a list of valid iSCSI targets (for example, volumes or snapshots) and making those targets available for use by the initiator.

failback

Mechanism for temporarily hosting a volume on the secondary group and then returning to the primary group. You may want to use failback if the primary group is unavailable due to failure or maintenance.

To perform failback, you can promote the volume's inbound replica set to a recovery volume (allowing initiators to access the recovery volume, as needed), demote the original volume (when it becomes available) to a failback replica set, replicate volume changes to synchronize the groups, demote the recovery volume to the original inbound replica set, and then promote the failback replica set to the original volume.

When configuring a volume for replication, if you enable the option to retain the failback snapshot, you may be able to fail back to the primary group by replicating only the changes that occurred while the volume was hosted on the secondary group.

failback baseline

Indicates the point at which you may be able to fail back to the primary group by replicating only the changes that occurred while the volume was hosted on the secondary group. A failback baseline depends on the existence of a failback snapshot. See *failback* on page A-2 and *failback snapshot* on page A-3.

failback replica set

Replica set created by demoting a replicated volume. To return to the original volume, you must promote the failback replica set. See *failback* on page A-2, *demote* on page A-2, and *promote* on page A-6.

failback snapshot

Snapshot that represents the most recent complete replica and establishes the failback baseline. When configuring a volume for replication, if you enable the option to retain the failback snapshot, you may be able to fail back to the primary group by replicating only the changes that occurred while the volume was hosted on the secondary group. See *failback baseline* on page A-2.

Gigabit Ethernet

Group of Ethernet standards in which data is transmitted at a rate of one Gigabit per second.

group

See *PS Series group* on page A-6.

group administrator

A user who has permission to manage all features and objects in a PS Series group, including configuring replication partners. Compare to *pool administrator*.

group IP address

Highly available IP address that is used by iSCSI initiators and administrators to access a group.

group member

See *member* on page A-5.

host

Any computer directly connected to a network that acts as a repository for services (such as e-mail or FTP) available to other computers on the network.

host bus adapter

I/O adapter that links an *iSCSI initiator* with target systems.

hot-swapping

Removing a component from a PS Series array and installing a replacement while the power is on, the array is running, and without disrupting I/O.

iSCSI

Acronym for Internet Small Computer Systems Interface. Standard protocol for transporting storage I/O block data over an IP network.

iSCSI initiator

Hardware or software that starts the transfer of information to and/or from an *iSCSI target* (for example, a volume or snapshot).

iSCSI target

SCSI block storage device that is accessed through TCP/IP and iSCSI protocols. Each volume and snapshot is seen on the network as an iSCSI target.

iSNS

Acronym for Internet Storage Name Service. Facilitates the automated discovery, management, and configuration of iSCSI and other devices on a TCP/IP network.

latency

Amount of time needed to complete a specific I/O operation. Also called *delay*.

manual transfer replication

The process of sending the first replica of a volume (the entire volume contents) to the partner location using physical media such as a CD-ROM, instead of the 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 replication applies only to the first data transfer for a volume.

Manual Transfer Utility

A stand-alone utility that enables you to perform the first volume replication using external media, instead of the network. The utility has two versions, a graphical user interface and an interactive command line interface.

member

A PS Series array that is connected to the network and has been configured into a PS Series group. Groups can have several members, optionally configured into one or more storage pools.

MIB

Acronym for Management Information Base (or Bases, when written as MIBs). A type of database used to manage the devices in a communications network. SNMP, a communication protocol between management stations (consoles, for example) and managed objects (such as routers, gateways, and switches) makes use of MIBs.

mirroring

Writing and maintaining the same (redundant) data in two equal-sized segments residing on two separate disks to protect data against loss due to disk or device I/O failure.

MRTG

Acronym for Multi Router Traffic Grapher, a third-party utility to monitor group performance.

multi-subnet configuration

Network configuration in which some of the IP addresses assigned to the network interfaces on the members and the group IP address reside on different subnets.

netmask

Network address used to identify the network segment (subnet) on which an IP address (network interface) resides. Same as subnet mask.

notification

Method a group uses to automatically inform users of significant events through e-mail, remote syslog files, and SNMP traps.

NTP

Acronym for Network Time Protocol. Defines a set of procedures for synchronizing clocks on hosts connected to a network and having access to the Internet.

parity

Form of redundancy used to re-create the data on a failed disk.

pool

Portion of group disk space to which you assign members and volumes. Pools provide you with greater control of group storage space and enable you to segregate users and applications, while maintaining a single-system view of the storage. Also called storage pool.

pool administrator

A user who has permission to manage objects only in a specific pool or set of pools for a group. Compare to *group administrator*.

primary volume

A volume configured for replication to a partner group.

promote

Converts an inbound replica set to a temporary recovery volume or converts a failback replica set to a volume, returning it to its original state. An inbound replica set can also be promoted to a volume that cannot be demoted (a permanent promotion). Also, see *demote* on page A-2 and *failback* on page A-2.

PS Series group

One or more PS Series arrays (members) configured on a network, accessed through a single IP address, and managed as a single system.

RADIUS

Acronym for Remote Authentication Dial-in User Service. This service centralizes the authentication and management (and optionally accounting) of user accounts and login information for multiple network connections.

RAID policy

The combined RAID level (RAID 10, RAID 5, or RAID 50) and spare disk policy for a member.

RAID

Acronym for Redundant Array of Independent Disks. The disks are joined together and managed as a single element. RAID arrays set aside part of the physical storage capacity to store redundant information about the data stored on the array. This information can then be used to regenerate user data in the event a disk in the array fails or the access to the disk fails.

RAID 0 (striping)

RAID level that distributes (stripes) data evenly across disks in equal-sized segments. Striping improves performance but does not provide fault tolerance.

RAID 1 (*mirroring*)

RAID level in which data is mirrored across (written to) two disk drives for redundancy and high availability.

RAID 5 (distributed parity)

RAID level in which parity information is distributed across all drives in the array for redundancy and high availability.

RAID 10

RAID level consisting of striping on top of multiple RAID 1 sets. Also called performance optimization.

RAID 50

RAID level consisting of striping on top of two RAID 5 sets. Also called capacity optimization.

primary group

In a replication operation, the group where the base (or primary) volume resides. See *secondary group* on page A-9.

recovery group

In a replication operation, the group (also called the secondary group) where volume data can be recovered if the primary group where the original volume resides is destroyed.

recovery volume

Temporary volume created by promoting an inbound replica set. To return to the original inbound replica set, you must demote the recovery volume. See *demote* on page A-2, *failback* on page A-2, and *promote* on page A-6.

replica

Represents the contents of a volume at a specific point in time, located on the secondary group in the replication operation. You can clone a replica or promote a replica set to recover volume data.

replica set

Set of complete replicas for a volume. You can promote a replica set to recover volume data.

replication

Functionality that enables you to replicate (copy) volumes from the primary group, where the volume resides, to the secondary group, where volume data can be recovered in the event of a disaster. The groups can be separated by an unlimited distance.

replica collection

One or more replicas created at the same time from a volume collection, containing one replica for each volume in the collection.

replica collection set

The set of replica collections for a volume collection.

reported size

Volume size as seen by iSCSI initiators. For a volume without thin provisioning enabled, the entire amount of space indicated by the reported size is allocated to the volume (that is, the reported size is the same as the volume reserve). For a volume with thin provisioning enabled, only the volume reserve is allocated to the volume. Therefore, the reported size may differ from the volume reserve, which increases as the thin-provisioned volume becomes in use. See *volume reserve* on page A-11.

SAN

Acronym for Storage Area Network, a high-speed subnetwork of shared storage devices. SAN architecture enables you to consolidate storage and simplify storage management. As more storage devices are added to a SAN, they too will be accessible from any server in the larger network. PS Series groups are iSCSI SANs, in which storage devices are accessible from an Ethernet network.

SAS

Acronym for Serial Attached SCSI technology. SAS is a serial communication protocol for direct attached storage (DAS) devices.

SATA

Acronym for Serial ATA (Advanced Technology Attachment) technology. SATA is the successor to the legacy ATA standard.

SCSI

Acronym for Small Computer System Interface.

secondary control module

In a dual control module array, the control module that mirrors data from the active control module.

secondary group

In a replication operation, the group (also called the recovery group) where the replicas are stored. Volume data can be recovered on the secondary group if the primary group, where the original volume resides, is destroyed.

SES

Acronym for SCSI Enclosure Services. A SCSI command set is defined for managing various non-SCSI elements contained within the enclosure.

single-subnet configuration

Network configuration in which the group IP address and the IP addresses assigned to the group member network interfaces reside on the same subnet.

SMART

Acronym for Self-Monitoring Analysis and Reporting Technology. This technology is designed to determine the reliability status of a disk. If the SMART system determines that a failure is imminent, the member would fail the disk and notify the user.

SMTP

Acronym for Simple Mail Transfer Protocol. SMTP is the protocol used for routing e-mail across the Internet.

snapshot

Represents the contents of a volume at a specific point in time and is seen on the network as an iSCSI target. Can be used for backups, testing, and upgrades. To recover data, you can restore a volume from a snapshot or clone a snapshot.

snapshot collection

One or more snapshots created at the same time from a volume collection or through a multi-volume snapshot operation. See *custom snapshot collection* on page A-2.

SNMP

Acronym for Simple Network Management Protocol.

spare (disk)

PS Series array disk used as a hot standby (redundant) component. A spare disk is powered on but does not operate until a data disk fails. Spare disks are configured and used automatically.

SSH

Acronym for Secure Shell, a program developed to log into another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels.

SSL

Acronym for Secure Socket Layer, a protocol developed by Netscape for transmitting private documents via the Internet. SSL works by using a public key to encrypt data that is transferred over the SSL connection. Many websites use SSL to obtain confidential user information, such as credit card numbers. By convention, URLs that require an SSL connection start with `https:` instead of `http:`.

storage pool

See *pool* on page A-6.

subnet mask

Network address used to identify the network segment (subnet) on which an IP address (network interface) resides. Same as netmask.

target portal

A target or iSCSI portal is the IP and TCP port number pair for a target. The default port number is 3260.

thin provisioning

Feature that allows administrators to improve storage resource utilization by logically over-allocating SAN volumes. With thin provisioning, you can meet current and future storage capacity needs while eliminating the difficulty of adding storage to applications or operating systems that do not support online expansion.

When first created, the volume reserve for a thin-provisioned volume is only a portion of the volume size. As the volume is used, the volume reserve will grow, up to a user-defined limit, at which point the volume is set offline. A warning is issued as additional space is allocated, enabling you to make adjustments, if desired.

throughput

Rate at which an I/O subsystem or component can perform I/O operations.

vacate

Operation in which all the volume data on a member is moved to a different member. For example, a member vacate operation occurs when you remove a member from the group or move a member across storage pools.

volume

Logical device representing a portion of storage pool space and seen on the network as an iSCSI target. Volume data is distributed across the pool members. Volumes are protected with access control records and can be expanded and load-balanced without disrupting access to the volume.

volume collection

One or more volumes grouped together for the purposes of creating snapshots or replicas at the same time.

volume reserve

Amount of pool space allocated to a volume (consumed from free pool space). For a volume without thin provisioning enabled, the volume reserve is the same as the reported size (seen by initiators). For a thin-provisioned volume, the volume reserve is initially set to a minimum value and then increases automatically (if space is available) as the volume becomes in use. See *reported size* on page A-8 and *thin provisioning* on page A-11.

B 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:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. 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.
3. 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.

4. Neither the name of the NetBSD Foundation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE NETBSD FOUNDATION, INC. AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE FOUNDATION OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT

LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This code is derived from software contributed to The NetBSD Foundation by Charles M. Hannum and by Jason R. Thorpe of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

This code is derived from software contributed to The NetBSD Foundation by John T. Kohl and Charles M. Hannum.

This code is derived from software contributed to The NetBSD Foundation by Kevin M. Lahey of the Numerical Aerospace Simulation Facility, NASA Ames Research Center.

This code is derived from software contributed to The NetBSD Foundation by Jun-ichiro Hagino.

This product includes software developed by Christopher G. Demetriou for the NetBSD Project.

Copyright © 1995, 1996, 1997, 1998 Christopher G. Demetriou.

This code is derived from software contributed to The NetBSD Foundation by Luke Mewburn.

This code is derived from software contributed to The NetBSD Foundation by Klaus Klein.

This code is derived from software contributed to The NetBSD Foundation by Jonathan Stone.

This code is derived from software contributed to The NetBSD Foundation by Jason R. Thorpe.

This code is derived from software contributed to The NetBSD Foundation by UCHIYAMA Yasushi.

This product includes software developed for the NetBSD Project by Wasabi Systems, Inc.

Copyright © 2000-2001 Wasabi Systems, Inc. All rights reserved.

This product includes software developed by the University of California, Berkeley and its contributors. This product includes software developed by the University of California, Lawrence Berkeley Laboratory.

Copyright 1985-1995 The Regents of the University of California.

Copyright 1997-2000 Niels Provos.

This code is derived from software contributed to Berkeley by Ralph Campbell.

This code is derived from software contributed to Berkeley by Rick Macklem.

Copyright © 1989 Digital Equipment Corporation.

This product includes software developed by Manuel Bouyer.

Copyright © 1999 Manuel Bouyer.

This product includes software developed by Adam Glass.

Copyright © 1994 Adam Glass.

This code is derived from software contributed to Berkeley by Paul Vixie.

This code is derived from software contributed to Berkeley by Chris Torek.

This code is derived from software contributed to Berkeley by Mike Hibler.

This code is derived from software contributed to Berkeley by Paul Borman at Krystal Technologies.

This code is derived from software contributed to Berkeley by Peter McIlroy.

This code is derived from software contributed to Berkeley by Peter McIlroy and by Dan Bernstein at New York University.

This code is derived from software contributed to Berkeley by Stephen Deering of Stanford University.

This code is derived from software contributed to Berkeley by Jeffrey Mogul.

Copyright 1996 The Board of Trustees of The Leland Stanford Junior University.

This product includes software developed by the Computer Systems Laboratory at the University of Utah. Copyright © 1990,1994 The University of Utah and the Computer Systems Laboratory (CSL). All rights reserved.

This code is derived from software contributed to Berkeley by the Systems Programming Group of the University of Utah Computer Science Department.

Copyright (c) 2000 Soren S. Jorvang.

Copyright (c) 1993 John Brezak. All rights reserved.

Copyright © 1995 - 2000 WIDE Project. All rights reserved.

© UNIX System Laboratories, Inc.

All or some portions of this file are derived from material licensed to the University of California by American Telephone and Telegraph Co. or Unix System Laboratories, Inc. and are reproduced herein with the permission of UNIX System Laboratories, Inc.

Copyright © 1999 Shuichiro URATA.

This product includes software developed by Matthias Pfaller.

Copyright © 1996 Matthias Pfaller.

Copyright © 1993 Jan-Simon Pendry.

This product includes software developed by Gordon W. Ross.

Copyright © 1995 Gordon W. Ross.

This product includes software developed by Philip A. Nelson.

Copyright © 1993 Philip A. Nelson.

Copyright © 1999 Ross Harvey.

This product includes software developed by Christos Zoulas.

Copyright © 1996 Christos Zoulas.

Copyright © 1997 Zubin D. Dittia.

This product includes software developed by SiByte, Inc.

Copyright © 2000 SiByte, Inc.

Copyright © 1996, 2000 Intel Corporation.

Copyright 1996 - 1998 Microsoft Corporation.

Copyright © 1990,1994 The University of Utah and the Computer Systems Laboratory (CSL).

Copyright © 1991 Bell Communications Research, Inc. (Bellcore).

Copyright © 2000 Caldera Systems, Inc. All rights reserved.

Copyright © 1995 - 2000 Kungliga Tekniska Högskolan.

(Royal Institute of Technology, Stockholm, Sweden). All rights reserved.

Copyright © 1993-1995 HEWLETT-PACKARD COMPANY

Copyright © 1995-1997 Eric Young All rights reserved.

Copyright © 1992 Simmule Turner and Rich Salz. All rights reserved.

Copyright © 1999 - 2001, PADL Software Pty Ltd. All rights reserved.

Copyright © 1985 - 1988 by the Massachusetts Institute of Technology.

Copyright © 1995 by Wietse Venema. All rights reserved.

Copyright © 1999 The OpenSSL Project. All rights reserved.

Copyright © 1992 – 1999 Theo de Raadt. All rights reserved.

Copyright © 1999 Dug Song. All rights reserved.

Copyright © 2000-2002 Markus Friedl. All rights reserved.

Copyright © 2001 Per Allansson. All rights reserved.

Copyright © 1998 CORE SDI S.A., Buenos Aires, Argentina.

Copyright © 2001-2002 Damien Miller. All rights reserved.

Copyright © 2001 Kevin Steves. All rights reserved.

Copyright © 1999 Aaron Campbell. All rights reserved.

Copyright © 2002 Nils Nordman. All rights reserved.

Copyright © 2000 Todd C. Miller. All rights reserved.

Copyright © 1995, 1996 by David Mazieres.

Copyright © 2000 Zembu Labs, Inc. All rights reserved.

Copyright © 2000 Takuya SHIOZAKI. All rights reserved.

Copyright © 1992 Keith Muller.

Copyright © 1994, Jason Downs. All rights reserved.

Copyright © 1997 Matthew R. Green. All rights reserved.

Copyright © 1999 Citrus Project. All rights reserved.

Copyright © 1990-2, RSA Data Security, Inc. All rights reserved.

Copyright © 1995 by International Business Machines, Inc.

Copyright © 1996 by Internet Software Consortium.

Copyright © 1995, 1999 Berkeley Software Design, Inc. All rights reserved.

Copyright © 1993 Carlos Leandro and Rui Salgueiro Dep. Matematica Universidade de Coimbra, Portugal, Europe.

Copyright © 1992, 1993, 1994 Henry Spencer.

Copyright © 1986-1991 by Sun Microsystems, Inc.

Copyright © 1993 Martin Birgmeier.

Copyright © 1991 by AT&T.

Copyright © 1997 Frank van der Linden. All rights reserved.

Copyright © 1999 Michael Graff. All rights reserved.

This product includes software developed by Alistair G. Crooks.

Copyright © 1999 Alistair G. Crooks. All rights reserved.

Copyright © 2001 Cerbero Associates Inc.

Copyright © 1995-1998 Mark Adler.

Copyright © 1995-1998 Jean-loup Gailly.

Copyright © 1998-1999 Brett Lymn. All rights reserved.

Copyright © 1996-1999 SciTech Software, Inc.

Copyright © 2001, 2002 Brian Stafford.

Copyright © 1999-2001 Bruno Haible.

Copyright © 2001 Alex Rozin, Optical Access. All rights reserved.

Copyright © 1989 TGV, Incorporated. All rights reserved.

Copyright © 2000 Frank Strauss. All rights reserved.

Copyright © 1997 Niels Baggesen. All rights reserved.

Copyright © 2000 National Aeronautics & Space Administration. All rights reserved.

Copyright © 1990-2000 Robin's Nest Software, Inc.

Copyright © 1989-1996 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms are permitted provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that the software was developed by Carnegie Mellon University. The

name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Dell is a trademark of Dell, Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Information in this document is subject to change without notice.

Reproduction in any manner whatsoever without the written permission is strictly forbidden.

Index

A

- access control records
 - asterisks in 5-26
 - creating 5-7, 5-24
 - deleting 5-27
 - displaying 5-26, 6-10
 - manual replication requirements
 - 8-6, 8-14, 8-21, 8-25
 - maximum for a volume 5-23
 - modifying 5-26
 - restricting host access to volumes and snapshots 5-23
- accounts (administration)
 - deleting local accounts 3-13
 - group administrator 3-9
 - grpadmin 3-9
 - managing 3-9
 - modifying local accounts 3-13
 - pool administrator 3-9
 - read-only 3-9
 - setting up local accounts 3-9
 - types 3-9
 - using RADIUS accounting 3-13, 3-18
 - using RADIUS authentication 3-13, 3-16
- accounts (local CHAP)
 - deleting 3-28
 - modifying 3-28
 - restricting iSCSI access 3-25
 - setting up 3-26
- accounts (RADIUS CHAP)
 - setting up 3-28
 - specifying RADIUS servers 3-28
- active control module
 - failover behavior 9-27
 - identifying 9-23
- alarms
 - displaying 10-4
 - format 10-5
 - GUI notification 2-14
 - priorities 10-4
- arrays
 - adding to group 2-2, 2-5
 - battery status 9-24
 - control module status 9-23
 - cooling modules 9-32
 - disk configuration 9-12
 - displaying hardware information 9-8
 - enclosure status 9-32
 - flashing LEDs 9-31
 - network configuration 9-18
 - power supplies 9-32
 - removing from group 9-30
 - See also members
 - serial connections 2-4
- authentication (administration
 - accounts)
 - using local accounts 3-9
 - using RADIUS servers 3-9, 3-13
- authentication (iSCSI)
 - initiators 3-25, 3-28, 5-7, 5-23
 - mutual 3-26, 3-29
 - targets 3-26, 3-29
 - using local CHAP 3-26
 - using RADIUS servers 3-26, 3-28

- authentication (replication)
 - between replication partners 7-2, 7-12
 - mutual 7-2

B

- backing up data
 - cloning replicas 7-42
 - cloning volumes 5-38
 - cloning, snapshots 6-15
 - promoting a replica set 8-30
 - replication 7-1, 8-1
 - snapshots 6-1
- batteries
 - setting policies 9-26
 - status 9-24

- binding a volume to a member 5-37

C

- cache mode policies, setting
 - low-battery-safe 9-26
 - single-controller-safe 9-26
- cache modes
 - write-back 9-25
 - write-through 9-25
- cancelling operations
 - member move 9-28
 - replication 7-41
 - volume move 5-29
- capacity
 - disks 9-16
 - displaying group 3-1
 - displaying member 9-6
 - displaying pool 4-6
 - increasing group 2-5
 - increasing member 9-16

- increasing pool 9-27

CHAP

- authenticating initiators and targets 3-25
- controlling iSCSI access 3-25
- creating local accounts 3-26
- initiator authentication 5-7, 5-23
- restricting access to volumes and snapshots 5-23
- target authentication 3-29
- using local accounts 3-25
- using RADIUS servers 3-28

CLI

- accessing group 3-7
- displaying events 10-4
- managing a group 2-19

cloning

- description 8-28
- replicas 7-42
- snapshots 6-15
- space requirements 5-38
- volumes 5-38

- collections, See volume collections, snapshot collections, and replica collections

connections (iSCSI)

- displaying member 9-35
- displaying snapshot 6-6
- displaying volume 5-11

control modules

- active 9-23
- batteries 9-26
- cache modes 9-25
- failover 9-23
- handling failures 9-27
- managing 9-23

- secondary 9-23
- status 9-23
- cooling modules
 - status 9-32
- custom snapshot collections
 - creating 6-18
 - defined 6-2
 - deleting 6-25
 - displaying 6-22
 - modifying 6-23
 - naming 6-20
- D**
- data recovery
 - procedures for 8-28
- date, setting in group 3-6
- default gateway, members 9-22
- delegated space
 - configuring 7-12
 - modifying 7-18
 - monitoring 7-33
 - planning 7-5
 - requirements 7-5, 7-12
 - sizing 7-5
- demoting a recovery volume to an inbound replica set 8-40
- demoting a volume to a failback replica set 8-47
- disabling firmware downgrades 3-40
- disaster recovery
 - cloning replicas 7-42
 - demoting a recovery volume 8-40
 - demoting a volume 8-47
 - failback 8-31

- methods 8-28
- promoting a failback replica set to a volume 8-48
- promoting a replica set permanently 8-44
- promoting a replica set temporarily 8-36
- using replication 7-1
- disks
 - adding to an array 9-16
 - capacity 9-16
 - failures 9-17
 - handling 9-16
 - installing 9-16
 - managing 9-16
 - replacing 9-16, 9-17
 - status 9-12, 9-13

E

- E-Mail Home notification, setting up 3-23
- e-mail notification, setting up 3-21
- events
 - displaying 10-1, 10-2, 10-4
 - E-Mail Home notification 3-20, 3-23
 - e-mail notification 3-20, 3-21, 10-4
 - GUI notification 2-14
 - informational 3-21, 3-24
 - message format 10-2
 - monitoring 10-1
 - notification methods 3-20
 - priorities 3-20, 10-1
 - syslog notification 3-20, 3-24, 10-4

F

failback

- demoting a recovery volume 8-40
- demoting a volume to a failback replica set 8-47
- description 8-31
- keeping failback snapshot 7-7, 7-21
- promoting a failback replica set 8-48
- promoting a replica set 8-36
- recovering from disasters 8-31
- steps 8-31

failback replica sets

- changing pools 8-50
- converting to inbound 8-51
- creating 8-47
- deleting 8-48, 8-51
- making demotion permanent 8-51
- pool change 8-50
- promoting to volume 8-48

failback replica space, displaying 7-39

failback snapshot

- description 7-21
- keeping 7-21
- using for failback 7-7

failover, control module 9-27

fans

- monitoring 10-15
- status 9-32

firmware

- disabling downgrades 3-40
- displaying 9-24
- member 9-7
- requirement for latest features 7-21

G

group

- access methods 3-7
- accounts (administration) 3-9
- accounts (CHAP) 3-26
- adding members 2-5
- CHAP authentication 3-25
- CLI access 3-7
- CLI management 2-19
- creating 2-2
- date 3-6
- description 1-2
- disk space 3-2
- event monitoring 10-1
- event notification 3-20
- events 10-1
- expanding capacity 2-5
- features 1-3
- getting started 2-1
- GUI access 3-7
- GUI management 2-7
- hardware installation 2-1
- informational messages 3-24
- IP address setting 2-3, 3-5
- iSCSI access 3-25
- iSNS servers 3-30
- load balancing 1-2, 3-38
- members 9-1
- MIBs 3-33
- monitoring 10-1
- naming 3-5
- network configuration 3-5
- network requirements 2-2
- NTP servers 3-6
- password for adding members 9-30
- pools 4-1
- removing members 9-30

- replication 7-1, 8-1
 - separating space 4-1
 - SMTP servers 3-21, 3-22, 3-23
 - snapshots 6-1
 - SNMP access 3-31
 - SNMP traps 3-32
 - status 3-1
 - steps for setting up 2-1
 - time 3-6
 - time zone 3-6
 - volume iSCSI settings, default 3-36
 - volume snapshot settings, default 3-36
 - volume thin-provisioning settings, default 3-36
 - volumes 5-1
- group defaults
- iSCSI target alias 3-38
 - snapshot reserve size 3-37
 - snapshot space recovery 3-37
 - thin-provisioning 3-38
- Group Manager
- online help 2-16
 - running GUI as a standalone application 2-17
 - using the CLI 2-19
 - using the GUI 2-7
- GUI
- accessing group 3-7
 - activities panel location 2-12
 - alarm notification 2-14
 - communication policies 2-13
 - controlling operation 2-11
 - data validation 2-15
 - debug logging 2-15
 - event log display 2-14
 - font size, controlling 2-12
 - help location 2-12, 2-16
 - installing as a standalone application 2-17
 - managing a group 2-7
 - popups 2-12
 - reconnecting 2-13
 - refresh data interval 2-13
 - requirements 2-7
 - running as a standalone application 2-17
 - secure access 2-11
 - tooltips 2-12
 - uninstalling standalone application 2-18
 - user preferences 2-11
 - window size recommendation 2-7
- ## H
- hardware
- installing 2-1
 - model information 9-36
 - status 9-32
- help, accessing 2-16
- hosts
- accessing volumes and snapshots 1-2
 - authenticating 3-25
 - displaying connections to members 9-35
 - displaying connections to volumes and snapshots 5-28
 - logging in to volumes and snapshots 5-27
 - TCP/IP port access to volumes 5-27

I

- iSCSI access
 - controlling 3-25
 - restricting target discovery 3-29
- iSCSI initiators
 - authentication 3-25, 5-7, 5-23
 - connecting to replica placeholder 8-25
 - connecting to replica snapshot 8-20
 - connecting to volumes and snapshots 5-27
 - description 1-2
 - iSNS discovery 3-30
 - restricting volume access with 5-25, 5-26
- iSCSI target alias
 - changing 3-38, 5-33
 - setting volume default 3-36
- iSCSI target name
 - copying 5-15
 - displaying for volumes 5-15
 - format for snapshots 6-2
 - format for volumes 5-2
- iSCSI targets
 - authenticating 3-26
 - description 1-2
 - restricting discovery 3-29
- iSNS, using for discovery 3-30, 5-27

L

- load balancing
 - disabling 3-38
 - enabling 3-38
 - types of 1-2, 1-6

- local replication reserve
 - handling insufficient space 7-9
 - keeping failback snapshot 7-7, 7-21
 - modifying 7-20
 - planning 7-7
 - sizing 7-7
 - specifying 7-20
 - using for failback 7-7
- lost blocks, handling in a volume 5-49

M

- manual transfer replication
 - confirming completion 8-1
 - description 8-1
 - enabling 7-21, 8-1
 - installing Manual Transfer Utility 8-2
 - transfer file size different from volume size 8-28
- Manual Transfer Utility
 - checking status using the GUI 8-17
 - confirming completion 8-16, 8-27
 - creating transfer files using the CLI 8-20, 8-24
 - creating transfer files using the GUI 8-4
 - description 8-1
 - file naming requirement 8-19
 - installing 8-2
 - loading transfer files using the CLI 8-25
 - loading transfer files using the GUI 8-11
 - steps for using CLI 8-19
 - steps for using the GUI 8-3

- transfer file size versus volume size 8-23
- members
 - adding to a group 2-5
 - assigning to a pool 4-2, 4-4, 9-28
 - battery status 9-24
 - binding a volume to 5-37
 - cancelling move operation 9-29
 - capacity, increasing 9-16
 - control modules
 - batteries 9-26
 - cache modes 9-25
 - handling failures 9-27
 - managing 9-23
 - status 9-23
 - cooling module status 9-32
 - default gateway 9-22
 - delay using disk space setting 9-5, 9-15
 - description 1-2
 - determining physical location 9-31
 - disabling downgrades 3-40
 - disks 9-16
 - adding 9-16
 - capacity 9-16
 - failures 9-17
 - handling 9-16
 - installing 9-16
 - removing 9-16
 - replacing 9-17
 - spares 9-12
 - status 9-12
 - displaying connections 9-35
 - displaying hardware information 9-8
 - displaying service information 9-36
 - flashing LEDs 9-31
 - health indicators 3-3
 - managing 9-1
 - monitoring hardware 9-34
 - network connectivity 9-18
 - network interfaces
 - configuring 9-19
 - deconfiguring 9-21
 - disabling 9-21
 - enabling 9-21
 - modifying 9-19
 - status 9-20
 - password for adding 9-30
 - pool assignment 4-7
 - power supply status 9-32
 - RAID policy 2-5, 2-7, 9-1, 9-3, 9-14
 - removing from group 9-30
 - renaming 9-29
 - status 9-6
 - synchronizing time with NTP 3-6
 - vacating 9-30
- MIBs, obtaining 3-33
- multi-subnet groups
 - connectivity between subnets 2-2
- mutual authentication
 - between initiators and targets 3-29
 - between replication partners 7-2, 7-12
- N**
 - network interfaces
 - configuring 9-19
 - deconfiguring 9-21
 - default gateway 9-22
 - disabling 9-21
 - enabling 9-21

- modifying 9-19
- requirements 9-18
- status 9-20
- networking
 - group IP address 3-5
 - recommendations 2-2
 - requirements 2-2, 9-18
 - subnet access 2-2
 - switch connectivity 2-2
- notification
 - e-mail 3-21
 - E-Mail Home 3-23
 - events 3-20
 - informational messages 3-24
 - syslog 3-24
- NTP, setting group time 3-6
- O**
- online help, accessing 2-16
- P**
- partners, see replication partners
- pools
 - administrator accounts 3-9
 - assigning members 9-28
 - assigning volumes 5-4, 5-29
 - changing for failback replica sets 8-50
 - characterizing application performance 4-4
 - characterizing member performance 4-4
 - creating 4-5
 - default 1-2, 4-2, 5-29
 - definition
 - deleting 4-11
 - displaying 4-6
 - failback replica set change 8-50
 - members in 4-7
 - merging 4-10
 - modifying description 4-10
 - modifying name 4-10
 - planning 4-4
 - segregating workloads 4-1
 - separating storage 4-2
 - space utilization 4-6
 - status 4-6
 - volumes in 4-8
- power supplies
 - status 9-32
- primary group, definition 7-2
- promoting a failback replica set to a volume 8-48
- promoting a replica set permanently 8-42, 8-44
- promoting a replica set to a recovery volume 8-36
- R**
- RADIUS servers
 - accounting for administration
 - accounts 3-13, 3-18
 - authenticating administration
 - accounts 3-13, 3-16
 - controlling iSCSI access 3-26
 - initiator authentication 3-28
- RAID
 - setting volume preference 5-35
 - supported RAID levels 9-2
- RAID policy
 - converting members 9-2, 9-14
 - description 9-1

- setting on a new member 2-5, 2-7, 9-3
- types of 9-1
- RAID-10 policy, definition 9-2
- RAID-10, no-spare policy, definition 9-2
- RAID-5 policy, definition 9-2
- RAID-50 policy, definition 9-2
- RAID-50, no-spare policy, definition 9-2
- recovering data
 - cloning a replica 7-42
 - cloning a snapshot 6-15
 - procedures 8-29
 - promoting a replica set
 - permanently 8-29, 8-44
 - promoting a replica set temporarily 8-29, 8-36
 - restoring a volume from a snapshot 6-14
- recovering data, procedures for 8-28
- recovery volume
 - creating 8-36
 - demoting to replica set 8-40
 - making promote permanent 8-42
 - restrictions 8-36
- Remote Setup Wizard, configuring a group 2-2, 2-5
- replica collection sets
 - defined 7-24, 7-31
 - displaying 7-31, 7-32
- replica collections
 - creating 7-23
 - defined 7-24, 7-31
 - deleting 7-47
 - displaying 7-31
 - enabling volume collection for replication 7-23
 - status 7-31
- replica reserve
 - description 7-20
 - handling insufficient space 7-7
 - modifying 7-20
 - planning 7-5
 - sizing 7-5
 - specifying 7-20
- replica sets
 - changing pools 8-50
 - deleting inbound 7-46
 - deleting outbound 7-45
 - displaying inbound 7-38
 - displaying outbound 7-35
 - naming 7-40
 - promoting permanently 8-29, 8-42, 8-44
 - promoting temporarily 8-29, 8-36
 - promoting to recovery volume 8-36
 - promoting to volume 8-48
- replicas
 - accessing data from 8-28
 - cancelling 7-41
 - cloning 7-42
 - creating from a schedule 7-25
 - creating from a volume 7-24
 - creating from a volume collection 7-24, 7-31
 - deleting inbound 7-46
 - deleting outbound 7-45
 - displaying inbound 7-38
 - displaying outbound 7-35

- manual transfer for 7-2
- scheduling 5-41, 7-25
- status 7-26, 7-29
- timestamp on each partner 7-36
- replication
 - authentication during 7-2
 - cancelling 7-41
 - communication between partners 7-12
 - configuration types 7-3
 - creating a replica 7-24
 - delegated space sizing 7-5
 - disabling 7-24
 - displaying partner activity 7-32
 - enabling 7-19
 - failback 8-31
 - failback replica space 7-39
 - failback snapshot 7-21
 - failback steps 8-31
 - impact of increasing volume reserve 7-11, 7-12
 - introduction 7-1
 - local replication reserve sizing 7-7
 - manual transfer 7-21, 8-1, 8-2
 - monitoring 7-29
 - delegated space 7-39
 - local replication reserve 7-30
 - operations 7-26
 - partners 7-32
 - replica collections 7-31
 - replica reserve 7-39
 - replicas 7-40
 - mutual authentication 7-9, 7-12
 - network requirements 7-9
 - partner space requirements 7-20
 - passwords for authentication 7-15
 - pausing 7-41
 - recovering data 8-28
 - replica reserve sizing 7-5
 - resuming 7-41
 - scheduling 5-41, 7-25
 - space requirements 7-5, 7-20
 - status 7-26, 7-27, 7-29
 - inbound 7-38
 - outbound 7-35
 - partners 7-32
 - steps for 7-9
 - stopping for a volume 7-24
 - volume configuration 7-19
- replication partners
 - configuring 7-11
 - definition 7-1
 - delegating space 7-12
 - deleting 7-18
 - displaying 7-32
 - displaying inbound replica sets 7-38
 - displaying outbound replica sets 7-35
 - firmware requirements 7-21
 - modifying 7-17
 - monitoring activity 7-32
 - monitoring delegated space 7-33
 - passwords for authentication 7-9, 7-12, 7-18
 - pausing replication between 7-41
 - requirements for 7-9
 - resuming replication between 7-41
 - status 7-32
- reptool, using for manual transfer replication 8-2

S

schedules

- disabling 5-47
- displaying 5-45
- enabling 5-47
- modifying 5-47
- replicas 5-41, 7-25
- snapshots 5-41, 6-14
- types 5-42
- volume collections 5-41

secondary control module

- failover behavior 9-27
- identifying 9-23

secondary group, definition 7-2

security

- administration accounts 3-9
- controlling access to group 3-7
- controlling access to volumes and snapshots 5-23
- requiring encryption for GUI 2-11

serial connection, instructions for 2-4

setup utility

- adding members to a group 2-5
- creating a group 2-2
- password for adding members 9-30

shrinking a volume. See volumes,
decreasing size.

SMTP, using for e-mail notification

- 3-21, 3-22, 3-23

snapshot collections

- creating 6-17
- definition 6-2
- deleting 6-25
- displaying 6-20

incomplete 6-21

modifying 6-23

naming 6-18

scheduling 6-16

status 6-20

snapshots

access control records

creating 5-7, 5-24

deleting 5-27

displaying 5-26, 6-10

modifying 5-26

cloning 6-15

controlling host access to 5-23

copy-on-write 6-3

creating from a volume 6-5

creating from a volume collection
6-17

default settings 3-36

definition 6-1

deleting automatically 6-3

deleting manually 6-16

displaying 6-6

host connections 5-28

iSCSI target alias 3-38, 6-12

iSCSI target name 6-2

logging in to 5-27

low-space warning threshold 3-37,
6-5

modifying settings 6-3

naming 6-2, 6-12

preserving 6-3

recovering space 6-5

reserving space 3-37, 6-4

restoring volumes from 6-14

scheduling 5-41, 6-14

setting offline 6-11

setting online 6-11

setting read-only 6-12

- setting read-write 6-12
 - space recovery policy 3-37, 6-5
 - status 6-6
 - SNMP
 - group access 3-31
 - trap destination 10-4
 - traps 3-32, 3-33
 - spares
 - displaying in a member 9-12
 - set as part of RAID policy 9-1
 - SSH, enabling for CLI access 3-7
 - status
 - batteries 9-24
 - control modules 9-23
 - cooling modules 9-32
 - custom snapshot collections 6-22
 - disks 9-13
 - free space 3-2
 - group 3-1
 - hardware 9-32
 - members 9-6
 - pools 4-6
 - power supplies 9-32
 - replica collections 7-31
 - replicas 7-26, 7-29
 - replication 7-26, 7-29
 - replication partners 7-32
 - snapshot collections 6-20
 - snapshots 6-6
 - volume collections 5-17
 - volumes 5-11
 - storage pools, See pools
 - syslog
 - event notification 3-20, 10-4
 - logging events 3-24
- T**
- TCP/IP port
 - volume access 5-27
 - telnet, enabling for CLI access 3-7
 - thin provisioning
 - automatic volume reserve growth 5-9
 - controlling space usage 5-9
 - description 5-8
 - disabling 5-34
 - enabling 5-5, 5-34
 - maximum in-use space 5-5
 - minimum volume reserve 5-5
 - restrictions 5-9
 - space issues 5-34
 - warning setting 5-5
 - time
 - setting manually 3-6
 - setting with NTP 3-6
 - time zone, setting 3-6
 - transfer files
 - creating 8-4
 - creating using the Manual Transfer Utility CLI 8-20
 - creating with Manual Transfer Utility GUI 8-4
 - loading using the Manual Transfer Utility CLI 8-25
 - loading using the Manual Transfer Utility GUI 8-11

V

vacating, members 9-30

VDS, Windows access to group 3-35

volume collections

- creating 5-15
- creating a replica collection 7-24
- creating a snapshot collection 6-17
- defined 5-15
- deleting 5-23
- disabling replication 7-24
- displaying 5-17
- enabling replication 7-23
- modifying 5-22
- scheduling replicas 5-41
- scheduling snapshots 5-41
- status 5-17

volumes

- access control records
 - creating 5-7, 5-24
 - deleting 5-27
 - displaying 5-26
 - modifying 5-26
- access controls 5-23
- assigning to a pool 4-2, 4-4, 5-4, 5-29
- backing up 6-1, 7-1, 8-1
- binding to a member 5-37
- cancelling move operation 5-30
- cloning 5-38
- collections 5-15
- controlling host access to 5-7, 5-23
- creating 5-3
- decreasing size 5-33
- default iSCSI settings 3-38
- default snapshot settings 3-37

default thin provisioning settings 3-38

deleting 5-49

demoting recovery volumes 8-40

demoting to failback replica sets 8-47

diagnosing problems 5-13

disabling all host access 5-23

disabling replication 7-24

disabling thin provisioning 5-34

displaying 5-11

enabling thin provisioning 5-5, 5-34

failback process 8-31

host connections 5-28

increasing size 5-31

iSCSI target alias 3-36, 3-38, 5-33

iSCSI target name 5-2, 5-14

keeping failback snapshot 7-21

load balancing 3-38

local replication reserve 7-7

logging in to 5-27

lost blocks 5-49

naming 5-3, 5-33

pool assignment 4-8

RAID preference 5-35

recovering data 8-28

replica reserve 7-5

replication 7-1, 8-1

- configuring 7-19

- manual transfer 7-21, 8-1, 8-2

- pausing 7-41

- requirements 7-9

- resuming 7-41

- setup steps 7-9

replication partners 7-11

reported size, description 5-2

- restoring from snapshot 6-14
- restricting access to 5-23
- restricting access with CHAP 3-25
- scheduling replication 5-41
- scheduling snapshots 5-41
- setting offline 5-48
- setting online 5-48
- setting read-only 5-48
- setting read-write 5-48
- shrinking 5-33
- snapshot reserve size 5-5, 6-4
- snapshot reserve warning
 - threshold 3-37, 6-5
- snapshot space recovery 6-5

- snapshots 6-1
- status 5-11
- TCP/IP port 5-27
- thin provisioning 5-8
- thin provisioning settings 5-5,
5-34
- volume reserve, description 5-2
- VSS, Windows access to group 3-35

W

- Web, enabling GUI access 3-7
- write-back mode 9-25
- write-through mode 9-25