

HP

P4000 SAN Solution User Guide

Abstract

This guide provides information for configuring and using the HP SAN Solution. It includes hardware configuration and information about designing and implementing a P4000 SAN. The intended audience is system administrators responsible for implementing, maintaining, and managing a P4000 SAN Solution.



© Copyright 2009, 2011 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgements

Microsoft, Windows, Windows XP, and Windows NT are U.S. registered trademarks of Microsoft Corporation.

Contents

1	Getting started.....	12
	Using the CMC.....	12
	Layout of the CMC.....	12
	Auto discover systems in the CMC.....	13
	Logging in.....	13
	Performing tasks in the CMC using the menu bar.....	13
	Configuring systems to add to management groups.....	14
	Using the Map View	14
	Using the display tools.....	14
	Using views and layouts.....	15
	Setting preferences.....	15
	Setting the font size and locale.....	15
	Setting naming conventions.....	15
	Creating storage by using the Getting Started Launch Pad.....	16
	Prerequisites.....	16
	Configuring storage systems.....	17
	Creating a volume using the wizard.....	17
	Enabling server access to volumes.....	18
	Finding storage systems after the first time.....	19
	Controlling which storage systems appear in the CMC.....	19
	Troubleshooting—Storage systems not found.....	19
	Setting up the CMC for remote support.....	20
2	Working with storage systems.....	21
	Storage system configuration categories.....	21
	Storage system configuration category definitions.....	21
	Storage system tasks.....	21
	Working with the storage system.....	21
	Logging in to and out of storage systems.....	22
	Changing the storage system hostname.....	22
	Locating the storage system in a rack.....	22
	Powering off or rebooting the storage system.....	23
	Powering on or off, or rebooting storage systems with modular components.....	23
	Rebooting the storage system.....	24
	Powering off the storage system.....	24
	Upgrading the SAN/iQ software on the storage system.....	25
	Upgrading the CMC and storage systems.....	25
	Setting upgrade preferences.....	25
	Checking for available upgrades.....	25
	Upgrading the CMC.....	26
	Upgrading storage systems in a management group or available storage systems.....	26
	Registering advanced features for a storage system.....	28
	Determining volume and snapshot availability.....	28
	Checking status of dedicated boot devices.....	29
	Checking boot device status.....	29
	Replacing a dedicated boot device.....	29
3	Storage Configuration: Disk RAID and Disk Management.....	30
	Configuring RAID and managing disks.....	30
	Getting there.....	30
	Status indicators.....	30
	Configuring and managing RAID.....	31

RAID Levels.....	31
Explaining RAID devices in the RAID setup report.....	32
RAID devices by RAID type.....	32
Planning the RAID configuration.....	33
Data protection.....	33
Using disk RAID with Network RAID in a cluster.....	34
Mixing RAID configurations.....	35
Setting RAID rebuild rate.....	35
General guidelines for setting the RAID rebuild rate.....	35
Setting the RAID rebuild rate.....	35
Reconfiguring RAID.....	36
To reconfigure RAID.....	36
Monitoring RAID status.....	36
Data reads and writes and RAID status.....	36
Data redundancy and RAID status.....	36
Managing disks.....	37
Getting there.....	38
Reading the disk report on the Disk Setup tab.....	38
Verifying disk status.....	39
Replacing a disk.....	42
Using Repair Storage System.....	43
Replacing disks in hot-swap storage systems.....	43
Preparing for a disk replacement.....	44
Replacing a disk in RAID 0.....	45
Replacing a disk in a hot-swap storage system	46
4 Managing the network.....	48
Network best practices.....	48
Changing network configurations.....	49
Managing settings on network interfaces.....	50
TCP status tab.....	50
Changing speed and duplex settings.....	50
Changing NIC frame size.....	51
Changing NIC flow control.....	52
The TCP/IP tab.....	53
Identifying the network interfaces.....	53
Pinging an IP address.....	54
Configuring the IP address manually.....	54
Using DHCP.....	55
Configuring network interface bonds.....	55
IP address for NIC bonds.....	56
NIC bonding and speed, duplex, frame size, and flow control settings.....	56
How Active-Passive bonding works.....	57
How link aggregation dynamic mode bonding works.....	60
How Adaptive Load Balancing works	62
Creating a NIC bond.....	64
Viewing the status of a NIC bond.....	67
Deleting a NIC bond.....	68
Disabling a network interface.....	70
Configuring a disabled interface.....	71
Using a DNS server.....	71
DNS and DHCP.....	71
DNS and static IP addresses.....	71
Adding the DNS domain name.....	71
Adding the DNS server.....	71

Adding domain names to the DNS suffixes.....	72
Editing a DNS server.....	72
Editing a domain name in the DNS suffixes list.....	72
Removing a DNS server.....	72
Removing a domain suffix from the DNS suffixes list.....	72
Setting up routing.....	73
Adding routing information.....	73
Editing routing information.....	73
Deleting routing information.....	73
Configuring storage system communication.....	74
Selecting the interface used by the SAN/iQ software.....	74
Updating the list of manager IP addresses.....	75
5 Setting the date and time.....	76
Management group time.....	76
Getting there.....	76
Refreshing the management group time.....	76
Using NTP.....	76
Editing NTP servers.....	77
Deleting an NTP server.....	77
Changing the order of NTP servers	77
Editing the date and time.....	78
Editing the time zone only.....	78
6 Administrative users and groups.....	79
Getting there.....	79
Managing administrative users.....	79
Default administrative user.....	79
Editing administrative users.....	79
Managing administrative groups.....	80
Default administrative groups.....	80
Editing administrative groups.....	81
7 Monitoring the SAN.....	83
Monitoring SAN status.....	83
Configuring the SAN Status Page	83
Using the SAN Status Page.....	84
Alarms and events overview.....	85
Working with alarms.....	87
Filtering the alarms list.....	87
Viewing and copying alarm details.....	87
Viewing alarms in a separate window.....	87
Exporting alarm data to a .csv file.....	88
Configuring events.....	88
Changing the event retention period.....	88
Setting up remote log destinations.....	88
Viewing events in a separate window.....	88
Working with events.....	89
Viewing new events.....	89
Filtering the events list.....	89
Viewing event details.....	90
Copying events to the clipboard.....	90
Exporting event data to a .csv or .txt file.....	90
Setting up email notification.....	91
Setting up the email server.....	91
Setting up email recipients.....	91

Setting up SNMP.....	92
Enabling SNMP agents.....	92
Adding SNMP traps.....	94
Using the SNMP MIBs.....	95
Running diagnostic reports.....	96
List of diagnostic tests.....	97
Generating a hardware information report.....	97
Saving a hardware information report.....	98
Hardware information report details.....	98
Using log files.....	100
Saving log files locally.....	100
Exporting the System Summary.....	100
Configuring a remote log and remote log destination.....	101
Editing remote log targets.....	101
Deleting remote logs.....	101
Exporting support logs.....	102
8 Working with management groups.....	103
Functions of management groups.....	103
Guide for management groups.....	103
Creating a management group.....	104
Creating a new management group.....	104
Best practice for managers in a management group.....	106
Managers overview.....	106
Functions of managers.....	106
Managers and quorum.....	106
Regular managers and specialized managers.....	107
Configuration Summary overview.....	108
Summary roll-up.....	108
Configuration guidance.....	109
Best Practice summary overview.....	111
Disk level data protection.....	112
Cluster-level data protection.....	112
Volume-level data protection.....	113
Volume access.....	113
Systems running managers.....	113
Network bonding.....	113
Network bond consistency.....	113
Network flow control consistency.....	113
Network frame size consistency.....	113
Management group maintenance tasks.....	113
Logging in to a management group.....	113
Logging out of a management group.....	114
Adding a storage system to an existing management group.....	114
Starting and stopping managers.....	114
Editing a management group.....	115
Saving management group configuration information.....	116
Safely shutting down a management group.....	116
Prerequisites.....	117
Start the management group back up.....	117
Removing a storage system from a management group.....	118
Prerequisites.....	118
Deleting a management group.....	119
Prerequisites.....	119
Setting the management group version.....	119

9 Using specialized managers.....	120
Failover Manager.....	120
Planning the virtual network configuration.....	120
Failover Manager requirements.....	120
Using the Failover Manager on Microsoft Hyper-V Server.....	121
Installing the Failover Manager for Hyper-V Server.....	121
Using the Failover Manager for VMware	122
Installing the Failover Manager for ESX Server	122
Installing the Failover Manager using the OVF files with the VI Client.....	123
Installing the Failover Manager for VMware Server or VMware Workstation.....	124
Troubleshooting the Failover Manager on ESX Server.....	124
Uninstalling the Failover Manager from VMware ESX Server.....	125
Virtual manager.....	125
When to use a virtual manager.....	125
Disaster recovery using a virtual manager.....	126
Storage system maintenance using a virtual manager.....	126
Requirements for using a virtual manager.....	126
Configuring a cluster for disaster recovery.....	127
Adding a virtual manager.....	129
Starting a virtual manager to regain quorum.....	130
Verifying virtual manager status.....	131
Stopping a virtual manager.....	131
Removing a virtual manager from a management group.....	131
10 Working with clusters.....	132
Clusters and storage systems.....	132
Creating a cluster.....	132
Cluster Map View.....	133
Monitoring cluster usage.....	133
Editing a cluster.....	133
Editing cluster properties.....	133
Editing iSNS servers.....	133
Editing cluster VIP addresses.....	134
Reconnecting volumes and applications after changing VIPs or iSNS servers.....	134
Maintaining storage systems in clusters.....	135
Adding a new storage system to a cluster.....	135
Upgrading the storage systems in a cluster using cluster swap.....	135
Reordering storage systems in a cluster.....	136
Exchange a storage system in a cluster.....	136
Removing a storage system from a cluster.....	136
Troubleshooting a cluster.....	136
Auto Performance Protection.....	136
Repairing a storage system.....	138
Deleting a cluster.....	139
11 Provisioning storage.....	140
Understanding how the capacity of the SAN is used.....	140
Provisioning storage.....	140
Provisioning volumes.....	140
Full provisioning.....	141
Thin provisioning.....	141
Planning data protection.....	141
Provisioning snapshots.....	146
Snapshots versus backups.....	146
The effect of snapshots on cluster space.....	146
Managing capacity using volume size and snapshots.....	147

Ongoing capacity management.....	147
Number of volumes and snapshots.....	147
Reviewing SAN capacity and usage.....	147
Measuring disk capacity and volume size.....	151
Changing the volume size on the server.....	152
Changing configuration characteristics to manage space.....	153
12 Using volumes.....	154
Volumes and server access.....	154
Prerequisites.....	154
Planning volumes.....	154
Planning how many volumes.....	154
Planning volume types.....	154
Guide for volumes.....	154
Creating a volume.....	156
Creating a basic volume.....	156
Configuring advanced volume settings [optional].....	157
Volumes map view.....	157
Editing a volume.....	157
To edit a volume.....	158
Deleting a volume.....	159
Restrictions on deleting volumes.....	159
Prerequisites.....	159
To delete the volume.....	160
13 Using snapshots.....	161
Types of snapshots.....	161
Uses and best practices for snapshots.....	161
Planning snapshots.....	162
Prerequisites for application-managed snapshots.....	163
Creating snapshots.....	163
Editing a snapshot.....	165
Scheduling snapshots.....	165
Best practices for scheduling snapshots of volumes.....	165
Requirements for snapshot schedules.....	166
Scheduling snapshots for volume sets.....	166
Creating a schedule to snapshot a volume.....	167
Mounting a snapshot.....	169
Mounting the snapshot on a host.....	169
Making a Windows application-managed snapshot available.....	170
Managing snapshot temporary space.....	172
Rolling back a volume to a snapshot or clone point.....	172
Rolling back a volume to a snapshot or clone point.....	173
Deleting a snapshot.....	175
14 SmartClone volumes.....	177
What are SmartClone volumes?.....	177
Prerequisites.....	177
SmartClone volume terminology.....	177
Example scenarios for using SmartClone volumes.....	178
Planning SmartClone volumes.....	179
Space requirements.....	179
Naming convention for SmartClone volumes.....	180
Server access.....	180
Defining SmartClone volume characteristics.....	180
Naming SmartClone volumes.....	181

Shared versus individual characteristics.....	182
Clone point.....	184
Shared snapshot	186
Creating SmartClone volumes.....	188
To create a SmartClone volume.....	188
Viewing SmartClone volumes.....	189
Map view.....	189
Editing SmartClone volumes.....	192
To edit the SmartClone volumes.....	193
Deleting SmartClone volumes.....	193
Deleting the clone point	193
Deleting multiple SmartClone volumes.....	194
15 Working with scripting.....	195
Scripting documentation.....	195
16 Controlling server access to volumes.....	196
Change in server access control from version 7.0 and earlier.....	196
Adding server connections to management groups.....	197
Prerequisites.....	197
Guide for servers.....	197
Adding a server connection.....	198
Managing server connections.....	198
Editing server connections.....	199
Deleting server connections.....	199
Clustering server connections.....	199
Requirements for clustering servers.....	199
Creating a server cluster.....	200
Server cluster map view.....	201
Working with a server cluster.....	201
Deleting a server cluster.....	201
Assigning server connections access to volumes.....	202
Assigning server connections from a volume.....	203
Assigning volumes from a server connection.....	203
Editing server connection and volume assignments.....	203
Completing the iSCSI Initiator and disk setup.....	204
Persistent targets or favorite targets.....	204
HP P4000 DSM for MPIO settings.....	204
Disk management.....	204
17 Monitoring performance	205
Prerequisites.....	205
Introduction to using performance information.....	205
What can I learn about my SAN?.....	205
Fault isolation example.....	206
What can I learn about my volumes?.....	207
Most active volumes examples.....	207
Activity generated by a specific server example.....	208
Planning for SAN improvements.....	208
Network utilization to determine if NIC bonding could improve performance example.....	208
Load comparison of two clusters example.....	209
Load comparison of two volumes example.....	209
Accessing and understanding the Performance Monitor window.....	210
Performance Monitor toolbar.....	211
Performance monitor graph.....	211
Performance monitor table.....	212

Understanding the performance statistics.....	213
Monitoring and comparing multiple clusters.....	215
Performance monitoring and analysis concepts.....	215
Workloads.....	215
Access type.....	215
Access size.....	216
Access pattern.....	216
Queue depth.....	216
Changing the sample interval and time zone.....	216
Adding statistics.....	216
Viewing statistic details.....	217
Removing and clearing statistics.....	218
Removing a statistic.....	218
Clearing the sample data.....	218
Clearing the display.....	218
Resetting defaults.....	218
Pausing and restarting monitoring.....	219
Changing the graph.....	219
Hiding and showing the graph.....	219
Displaying or hiding a line.....	219
Changing the color or style of a line.....	219
Highlighting a line.....	219
Changing the scaling factor.....	220
Exporting data.....	220
Exporting statistics to a CSV file.....	220
Saving the graph to an image file.....	221
18 Registering advanced features.....	222
Evaluation period for using advanced features.....	222
Starting the evaluation period.....	222
Backing out of Remote Copy evaluation.....	223
Scripting evaluation.....	223
Turn on scripting evaluation.....	223
Turn off scripting evaluation.....	224
Registering advanced features.....	224
Using license keys.....	224
Registering available storage systems for license keys.....	224
Registering storage systems in a management group.....	225
Saving and editing your customer information.....	227
19 iSCSI and the HP P4000 SAN Solution.....	229
Number of iSCSI sessions.....	229
Virtual IP addresses.....	229
Requirements for using a virtual IP address.....	229
iSNS server.....	229
iSCSI load balancing.....	230
Requirements.....	230
Authentication (CHAP).....	230
Requirements for configuring CHAP.....	231
iSCSI and CHAP terminology.....	231
Sample iSCSI configurations.....	232
Best practice.....	234
About HP DSM for MPIO.....	234
20 Using the Configuration Interface.....	235
Connecting to the Configuration Interface.....	235

Establishing a terminal emulation session on a Windows system.....	235
Establishing a terminal emulation session on a Linux/UNIX system.....	235
Opening the Configuration Interface from the terminal emulation session.....	236
Logging in to the Configuration Interface.....	236
Configuring administrative users.....	236
Configuring a network connection.....	236
Deleting a NIC bond.....	237
Setting the TCP speed, duplex, and frame size.....	237
Removing a storage system from a management group.....	238
Resetting the storage system to factory defaults.....	238
21 Support and other resources.....	239
Contacting HP.....	239
Subscription service.....	239
HP Insight Remote Support Software.....	239
New and changed information in this edition.....	240
Related information.....	240
HP websites.....	240
Customer self repair.....	240
A Replacing disks reference.....	241
Replacing disks and rebuilding data.....	241
Before you begin	241
Prerequisites.....	241
Replacing disks.....	242
Verify storage system not running a manager.....	242
Repair the storage system.....	242
Replace the disk.....	243
In the DL320s (NSM 2120), HP LeftHand P4300, HP LeftHand P4500.....	243
Rebuilding data.....	243
Recreate the RAID array.....	243
Checking progress for RAID array to rebuild.....	243
Return storage system to cluster.....	244
Restarting a manager.....	244
Add repaired system to cluster.....	245
Rebuild volume data.....	245
Controlling server access.....	245
Remove ghost storage system.....	246
Finishing up.....	246
B Third-party licenses.....	247
C SANiQ TCP and UDP Port Usage.....	248
Glossary.....	251
Index.....	257

1 Getting started

Welcome to the SAN/iQ software and the Centralized Management Console (CMC). Use the CMC to configure and manage the HP P4000 SAN Solution.

This product guide provides instructions for configuring individual storage systems, as well as for creating storage clusters, volumes, snapshots, and remote copies.

Using the CMC


Use the CMC to:

- Configure and manage storage systems
- Create and manage the P4000 SAN solution

Layout of the CMC

The CMC is divided into sections, as shown in [Figure 1](#) (page 12).

Figure 1 Parts of the CMC



1. Navigation window
2. Tab window
3. Alarms window
4. Status bar

Navigation window

The left vertical pane displays the architecture of your P4000 SAN. Access the Configuration Summary and Best Practice Summary in the navigation window, as well management groups, clusters and storage volumes.

- Many items have right-click menus.
- Each item in the navigation window has an icon depicting what type of item it is. A faded-looking icon indicates a remote item. Find a description of all the icons used in the CMC in **Help**→**Graphical Legend**.

Tab window

Tab windows display details about items selected in the navigation window, and they provide access to functions related to those items. For example, [Figure 1 \(page 12\)](#) shows the tabs that appear when a management group is selected in the navigation window.

Commands related to the selected items are accessible from the Tasks menu on the bottom left of the tab window.

NOTE: If you change the default size of the CMC application on your screen, the blue Tasks button at the bottom left of the tab window may be hidden. Scroll the tab window with the scroll bar to bring the Task button back into view.

Alarms window

Warning and critical events appear in the Alarms window for all management groups you are logged in to. Alarms stay in the Alarms window until the situation is resolved. For more information, see [“Alarms and events overview” \(page 85\)](#).

Auto discover systems in the CMC

When you open the CMC, it automatically searches the subnet for storage systems. Any storage systems that are found on the subnet appear in the navigation window on the left side of the CMC. If no storage systems are found automatically, the Find Systems Wizard opens and takes you through the steps to discover the storage systems on your network. For information about controlling which systems automatically appear when you open the CMC, see [“Finding storage systems after the first time” \(page 19\)](#).

Logging in

The CMC automatically logs in to storage systems in the Available Systems pool to access the system configuration categories. After you have created management groups, you must manually log in the management group. After you have logged in to one management group, the CMC attempts to log in automatically to other management groups using the first login.

Δ CAUTION: Do not log in to the same management group from more than one CMC.

Performing tasks in the CMC using the menu bar

The menu bar provides access to the following task menus:

- **File**—Lets you exit the CMC gracefully.
- **Find**—Finds storage systems that can be managed through the CMC.

- **Tasks**—Lets you access all storage configuration tasks. The tasks in this menu are grouped by logical or physical items. Tasks are also accessible through right-click menus and from the Tasks button in the tab window.
- **Help**—Lets you open the online help, access the online upgrades feature, and set CMC preferences. The Graphical Legend is a key to the icons used in the CMC and is available from the Help menu.

Configuring systems to add to management groups

Systems, including Failover Managers, that have not been added to management groups appear in the Available Systems. These systems are available to be added to management groups.

Other information in the navigation window depicts the storage architecture you create on your system. An example setup is shown in [Figure 1 \(page 12\)](#).

Using the Map View







The Map View tab is available for viewing the relationships between management groups, servers, sites, clusters, volumes and snapshots. When you log in to a management group, there is a Map View tab for each of those elements in the management group. For example, when you want to make changes such as moving a volume to a different cluster, or deleting shared snapshots, the Map View allows you to easily identify how many snapshots and volumes are affected by such changes.

The Map View pane contains display tools to control and manipulate the view. The display tools are available from the Map View Tasks menu or from the tool bar across the top of the pane. The tools function the same from either the tool bar or the Map View tasks menu.

Using the display tools

Use these tools, described in [Table 1 \(page 14\)](#), to select specific areas of the map to view, zoom in on, rotate, and move around the window. If you have a complex configuration, use the Map View tools to easily view and monitor the configuration.

Table 1 Map View display tools

Tool Icon	Function
	Zoom In—incrementally magnifies the Map View window.
	Zoom Out—incrementally reduces the Map View window.
	Magnify—creates magnification area, like a magnifying glass, that you can move over sections of the map view. Note that the magnify tool toggles on and off. You must click the icon to use it, and you must click the icon to turn it off.
	Zoom to Fit—returns the map view to its default size and view.
	Select to Zoom—allows you to select an area of the map view and zoom in on just that area.
	Rotate—turns the map view 90 degrees at a time.
Click and drag	You can left-click in the Map View window and drag the map around the window

Using views and layouts

The views and layouts differ for each element of your network that uses the map view.

For views and layouts available, see:

- Working with management groups, [“Management group map view tab”](#) (page 105)
- Controlling server access to volumes, [“Server cluster map view”](#) (page 201)
- Sites, *HP P4000 Multi-Site HA/DR Solution Pack User Guide*
- Clusters, [“Cluster Map View”](#) (page 133)
- Volumes and Snapshots, [“Volumes map view”](#) (page 157)
- SmartClone volumes, [“Using views”](#) (page 190)

Setting preferences

Use the Preferences window to set the following:

- Font size in the CMC
- Locale for the CMC. The locale determines the language displayed in the CMC.
- Naming conventions for storage elements
- Online upgrade options. See [“Setting upgrade preferences”](#) (page 25).

Setting the font size and locale

Use the Preferences window, opened from the Help menu, to set font size and locale in the CMC. Font sizes from 9 through 16 are available.

The CMC obtains the locale setting from your computer. If you change the locale on your computer and open the CMC, it uses the new locale, if available in the CMC.

You can override the locale setting from your computer by selecting a different locale in the CMC. Changing the locale in the CMC affects the language of the text that appears in the CMC and in the online help.

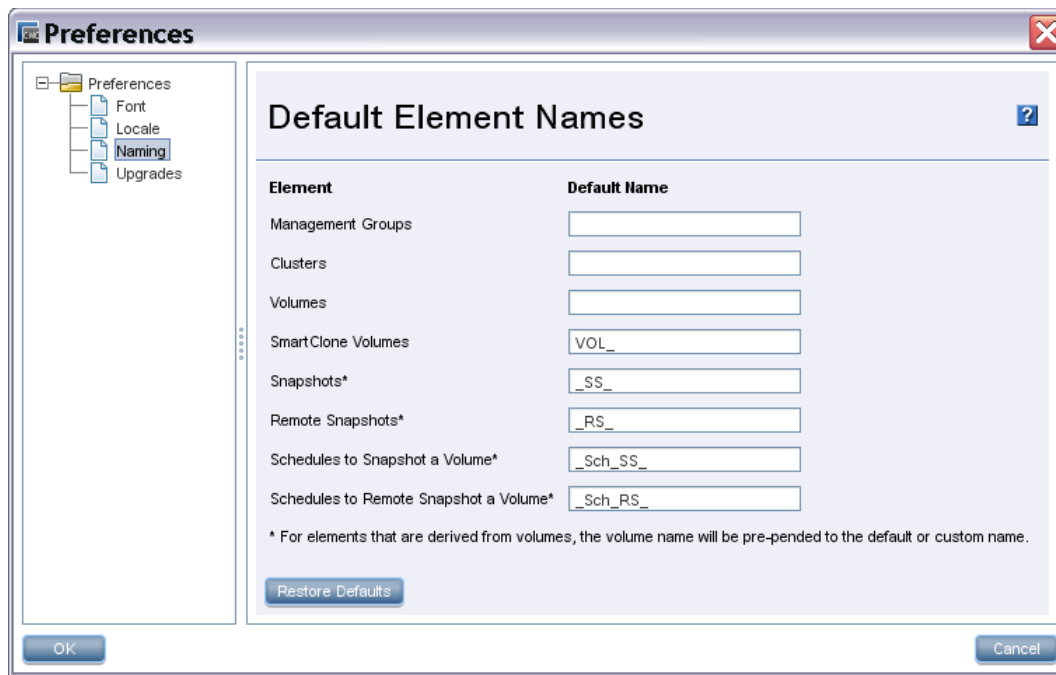
After changing the locale, you must close and reopen the CMC to see the change.

Setting naming conventions

Use the Preferences window, opened from the Help menu, to set naming conventions for elements you create when building the HP P4000 SAN Solution. You can use the default values or create your own set of customized values.

When you install the CMC for the first time, or upgrade from release 7.0.x, default names are enabled for snapshots, including schedules to snapshot a volume, and for SmartClone volumes. No default names are provided for management groups, clusters, and volumes.

Figure 2 Default naming conventions for snapshots and SmartClone volumes



Changing naming conventions

Change the elements that use a default naming convention or change the naming convention itself. If you use the given defaults, the resulting names look like those in [Table 2 \(page 16\)](#). Notice that the volume name carries into all the snapshot elements, including SmartClone volumes, which are created from a snapshot.

Table 2 Example of how default names work

Element	Default name	Example
SmartClone Volumes	VOL_	VOL_VOL_ExchLogs_SS_3_1
Snapshots	_SS_	VOL_ExchLogs_SS_1
Remote Snapshots	_RS_	VOL_RemoteBackup_RS_1
Schedules to Snapshot a Volume	_Sch_SS_	VOL_ExchLogs_Sch_SS_2.1
Schedules to Remote Snapshot a Volume	_Sch_RS_	VOL_ExchLogs_Sch_RS_2_Pri.1, VOL_RemoteBackup_Sch_RS_1_Rmt.1

If you delete all the default names from the Preferences Naming window, the only automatically generated naming elements that remain will incrementally number a series of snapshots or SmartClone volumes.

Creating storage by using the Getting Started Launch Pad

Follow the steps in this section to set up a volume quickly. Using the wizards on the Getting Started Launch Pad, work through these steps with one storage system, and with one strategy. The rest of this product guide describes other methods to create storage, as well as detailed information on features of the iSCSI SAN.

Prerequisites

- Install the storage systems on your network.
- Know the IP address you configured with the KVM or serial Configuration Interface when you installed the storage system.

- Install the HP P4000 CMC software on a management workstation or server that can connect to the storage systems on the network.
- Install an iSCSI initiator, such as the latest version of the Microsoft iSCSI Initiator, on the application server(s).

Finding storage systems

When you open the CMC, it searches for systems using Auto Discover by Broadcast. If no systems are found, the Find Systems window opens. Add individual IP addresses to find systems.

Configuring storage systems

Configure the storage system next. If you plan to use multiple storage systems, they must all be configured before you use them for clustered storage.

The most important categories to configure are:

- **RAID**—The storage system is shipped with RAID already configured and operational. Find instructions for changing RAID, and for ensuring that drives in the storage system are properly configured and operating in [“Storage Configuration: Disk RAID and Disk Management”](#) (page 30).
- **TCP/IP Network**—Bond the NIC interfaces and set the frame size, NIC flow control, and speed and duplex settings. Read detailed network configuration instructions in [“Managing the network”](#) (page 48).

To configure storage systems

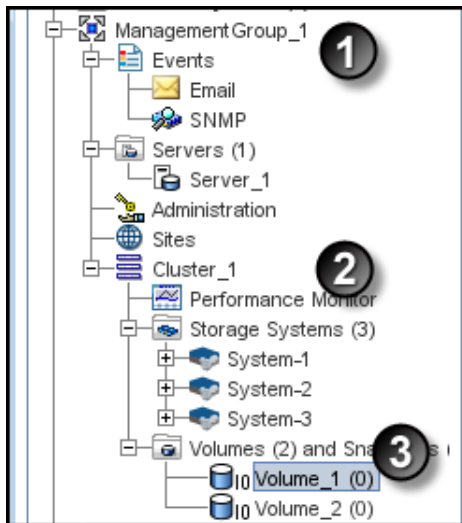
1. From the navigation window, select a storage system in the Available Systems pool.
2. Open the tree underneath the storage system.
3. In the list of configuration categories, select the Storage category.
4. Select the **RAID Setup** tab and verify the RAID settings or change the RAID level.
5. In the list of configuration categories, select the **TCP/IP Network** category and configure the network settings.

Creating a volume using the wizard

Next, you create storage volumes using the Management Groups, Clusters, and Volumes wizard, found on the Getting Started Launch Pad. Select **Getting Started** in the navigation window to access the Getting Started Launch Pad. On the Launch Pad, select the **Management Groups, Clusters, and Volumes Wizard**.

The wizard takes you through creating the tasks of creating a management group, a cluster, and a storage volume. This storage hierarchy is depicted in [Figure 3](#) (page 18).

Figure 3 The SAN/iQ software storage hierarchy



1. Management group
2. Cluster
3. Volume

To complete this wizard, you will need the following information:

- A name for the management group.
- A storage system discovered on the network and then configured for RAID and the TCP/IP Network settings
- DNS domain name, suffix, and server IP address for email event notification
- IP address or hostname and port of your email (SMTP) server for event notification
- A name for the cluster
- Virtual IP address to use for the cluster
- A name for the volume
- The size of the volume

NOTE: Names of management groups, clusters, volumes, and snapshots cannot be changed in the future without destroying the management group.

Enabling server access to volumes

Use the Assign Volume and Snapshot wizard to prepare the volume for server access. You set up application servers in the management group, then assign volumes to the servers. See [“Controlling server access to volumes” \(page 196\)](#) for a complete discussion of these functions.

To work through the Assign Volume and Snapshot wizard, you must first have created a management group, cluster, and at least one volume. You should also plan the following:

- The application servers that need access to volumes.
- The iSCSI initiator you plan to use. You need the server’s initiator name, and CHAP information if you plan to use CHAP.

Finding storage systems after the first time

The Find settings from your first search are saved in the CMC. Every time you open the CMC, the same search automatically takes place, and the navigation window is populated with all the storage systems that are found.

Control the systems that appear in the CMC based on criteria you set in the Find Systems window, opened from either the menu bar or the Getting Started Launch Pad. Choose to display a small group of storage systems, such as those in one management group, or display all storage systems on the subnet at one time.

Turn off Auto Discover by Broadcast for storage systems

If you do not want the CMC to automatically discover all the storage systems on the network when it opens, turn off Auto Discover by Broadcast.

1. From the menu bar, select **Find**→**Find Systems**.
2. Clear the **Auto Discover by Broadcast** check box.

The next time you open the CMC, it will not search the network for all storage systems.

Clearing the found storage systems from the navigation window

1. From the menu bar, select **Find**→**Clear All Found Items** to remove all storage systems from view in the navigation window.

Controlling which storage systems appear in the CMC

Control which storage systems appear in the navigation window by entering only specific IPs in the Find Systems window. Then, when you open the CMC, only those storage systems will appear in the navigation window. Use this method to control which management groups appear.

To control which storage systems appear

1. From the menu bar, select **Find**→**Find Systems**.
2. Click **Add** to enter the IP address of a storage system. Repeat for all the IP addresses of the desired set of storage systems.
3. Click **Find** to search for the list of IP addresses.

Troubleshooting—Storage systems not found

If the network has a lot of traffic, or if a storage system is busy reading or writing data, it may not be found when a search is performed. Try the following steps to find the storage system.

1. If the storage system you are looking for does not appear in the navigation window, search again using the Find menu.
2. If you have searched using Auto Discover by Broadcast, try adding individual IP Addresses and clicking **Find**.
3. If you have searched by Individual IP addresses, try searching by Auto Discover instead.
4. If searching again does not work, try the following:
 - Check the physical connection of the storage system.
 - Wait a few minutes and try the search again. If activity to the storage system was frequent, the storage system might not have responded to the search.

Possible reasons for not finding storage systems

Other problems can prevent the CMC from finding a storage system:

- Extremely high network traffic to and from the storage system.
- The IP address could have changed if the storage system is configured to use DHCP (not recommended).
- The storage system may have been rebooted and is not yet online.
- Power could have failed to a network switch that the storage system is connected to.
- The CMC might be running on a system that is on a different physical network than the storage system. Poor network routing performance at the site may severely affect performance of the CMC.

Setting up the CMC for remote support

If you are using HP remote support, the table below lists the required CMC setup.

Table 3 CMC setup for remote support

Requirement	For more information, see
SNMP enabled on each storage system	“Enabling SNMP agents” (page 92)
SNMP trap recipient set to IP address of the system where the remote support client is installed	“Adding SNMP traps” (page 94)
Port 8959 (used for the CLI) open	Your network administrator
Management group login and password for a read-only (View_Only_Administrator group) user	“Adding a new administrative user” (page 79)

2 Working with storage systems

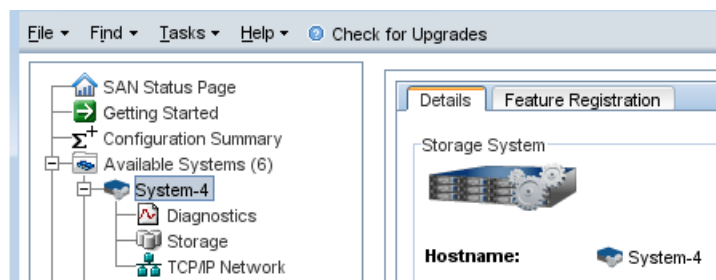
Storage systems displayed in the navigation window have a tree structure of configuration categories under them. The storage system configuration categories include:

- Diagnostics
- Storage
- TCP/IP Network

Storage system configuration categories

Storage system configuration categories allow access to all the configuration tasks for individual storage systems. You must log in to each storage system individually to configure, modify, or monitor the functions of that storage system.

Figure 4 Storage system configuration categories



Storage system configuration category definitions

- **Diagnostics**—Use the hardware category to run hardware diagnostic tests, to view current hardware status and configuration information, and to save log files.
- **Storage**—Manage RAID and the individual disks in the storage system.
- **TCP/IP Network**—For each storage system, configure and manage the network settings, including network interface cards (NICs), the routing table, and which interface carries SAN/iQ communication.

Storage system tasks

This section describes how to perform basic storage system tasks:

- “Working with the storage system” (page 21)
- “Logging in to and out of storage systems” (page 22)
- “Changing the storage system hostname” (page 22)
- “Locating the storage system in a rack” (page 22)
- “Rebooting the storage system” (page 24)
- “Powering off the storage system” (page 24)

Working with the storage system

After finding all the storage systems on the network, you configure each storage system individually.

1. Select the storage system in the navigation window.
Usually you will be logged in automatically. However, you will have to log in manually for any storage systems running a software version earlier than release 7.0. If you do need to manually log in, the Log In window opens.
2. Enter a user name and password.
3. Click **Log In**.

Logging in to and out of storage systems

You must log in to a management group to perform any tasks in that group. Logging into the management group automatically logs you into the storage systems in that group. You can log out of individual storage systems in the management group, and log back in to them individually.

Automatic login

Once you have logged in to a management group, additional log ins are automatic if the same user names and passwords are assigned. If management groups have different user names or passwords, then the automatic log in fails. In that case you must log in manually.

1. Enter the correct user name and password.
2. Click **Log In**.

Logging out of a storage system

1. Select a storage system in the navigation window.
2. Right-click, and select **Log Out**.

NOTE: If you are logged in to multiple storage systems, you must log out of each storage system individually.

Changing the storage system hostname

The storage system arrives configured with a default hostname. Use these steps to change the hostname of a storage system.

1. In the navigation window, log in to the storage system.
2. On the Details tab, click **Storage System Tasks** and select **Edit Hostname**.
3. Enter the new name, and click **OK**.
4. Click **OK**.

NOTE: Add the hostname and IP pair to the hostname resolution methodology employed in your environment, for example, DNS or WINS.

Locating the storage system in a rack

The Set ID LED turns on lights on the physical storage system so that you can physically locate that storage system in a rack.

NOTE: The Set ID LED is available depending on the storage system.

1. Select a storage system in the navigation window and log in.
2. Click **Storage System Tasks** on the Details tab and select **Set ID LED On**.
The ID LED on the front of the storage system is now a bright blue. Another ID LED is located on the back of the storage system.
When you click **Set ID LED On**, the status changes to On.
3. Select **Set ID LED Off** when you have finished.
The LED on the storage system turns off.

Powering off or rebooting the storage system

You can reboot or power off the storage system from the CMC. You can also set the amount of time before the process begins, to ensure that all activity to the storage system has stopped.

Powering off the storage system through the CMC physically powers it off. The CMC controls the power down process so that data is protected.

Powering off an individual storage system is appropriate for servicing or moving that storage system. However, if you want to shut down more than one storage system in a management group, you should shut down the management group instead of individually powering off the storage systems in that group. See [“Safely shutting down a management group”](#) (page 116).

Powering on or off, or rebooting storage systems with modular components

Some storage systems are comprised of modular components, that may include:

- Disk enclosure
- Server blades enclosure
- System controller

The P4800 is an example of a storage system that is comprised of modular components.

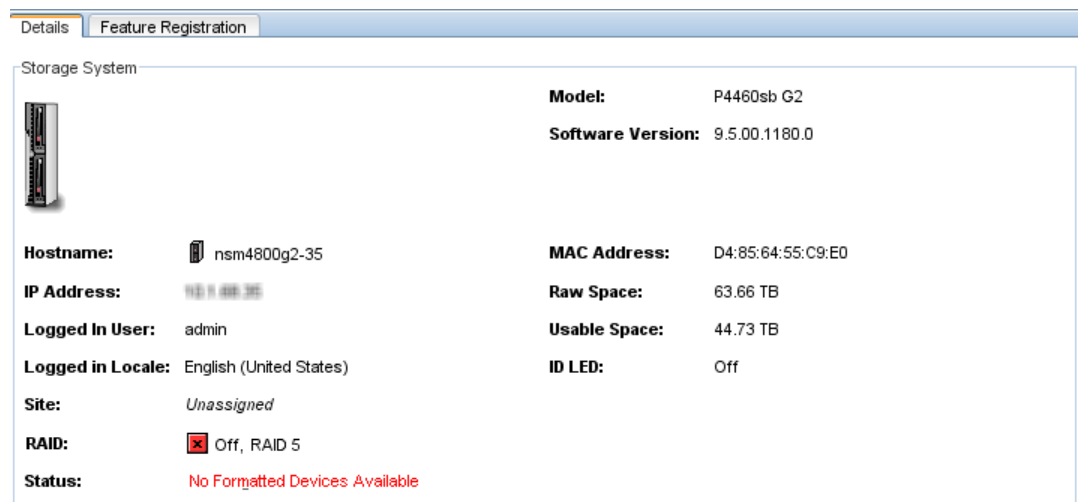
Power on the components in the following order:

1. Disk enclosure.
2. Server blades enclosure or system controller.

Allow up to 6 minutes for the system controller to come up completely and be discovered by the CMC. If you cannot discover the storage system using the CMC after 6 minutes, contact Customer Support.

3. If you do not power on the disk enclosure first, the Storage System Details tab shows the status with No Formatted Devices Available.

Figure 5 Disk enclosure not found as shown in Details tab



When powering off the storage system, be sure to power off the components in the following order:

1. Power off the server blades enclosure or system controller from the CMC as described in [“Powering off the storage system”](#) (page 24).
2. Manually power off the disk enclosure.

When you reboot the storage system, use the CMC, as described in [“Rebooting the storage system”](#) (page 24). This process reboots only the server blades enclosure or the system controller.

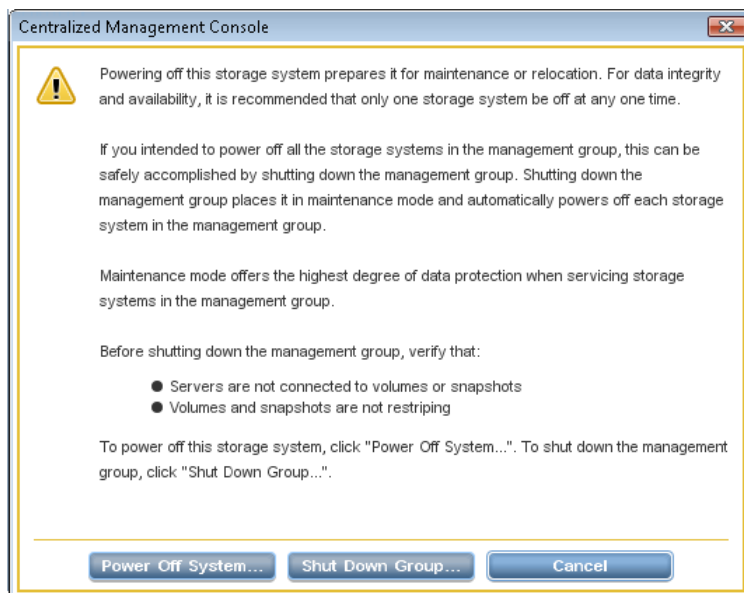
Rebooting the storage system

1. Select a storage system in the navigation window and log in.
 2. Click **Storage System Tasks** on the Details tab and select **Power Off** or **Reboot**.
 3. In the minutes field, enter the number of minutes before the reboot should begin.
Enter any whole number greater than or equal to 0. If you enter 0, the storage system reboots shortly after you confirm the reboot.
-
- NOTE:** If you enter 0 for the value when rebooting, you cannot cancel the action. Any value greater than 0 allows you to cancel before the reboot actually takes place.
-
4. Select **Reboot** to perform a software reboot without a power cycle.
 5. Click **OK**.
The storage system starts the reboot in the specified number of minutes. The reboot takes several minutes.
 6. Search for the storage system to reconnect the CMC to the storage system once it has finished rebooting.
See ["Finding storage systems"](#) (page 17).

Powering off the storage system

1. Log in to the storage system.
 2. Select **Storage System Tasks** on the Details tab and select **Power Off** or **Reboot**.
 3. Select **Power Off**.
The button changes to Power Off.
 4. In the minutes field, enter the number of minutes before the powering off should begin.
Enter any whole number greater than or equal to 0. If you enter 0, the storage system powers off shortly after you confirm the power off.
-
- NOTE:** If you enter 0 for the value when powering off, you cannot cancel the action. Any value greater than 0 allows you to cancel before the power off actually takes place.
-
5. Click **Power Off**.

Figure 6 Confirming storage system power off



Depending on the configuration of the management group and volumes, your volumes and snapshots can remain available.

Upgrading the SAN/iQ software on the storage system

The Upgrades tab shows what upgrades are available for your CMC and storage systems. Upgrades are available when the button on the menu bar says **Upgrades Available**. If the button says **Check for Upgrades**, click to see if upgrades are available.

When you upgrade the SAN/iQ software on a storage system, the version number changes. Check the current software version by selecting a storage system in the navigation window and viewing the Details tab window.

Upgrading the CMC and storage systems

For CMC and SAN/iQ upgrades, you can do the following:

- Set upgrade preferences. See [“Setting upgrade preferences”](#) (page 25).
- Check for available upgrades. See [“Checking for available upgrades”](#) (page 25).
- Upgrade the CMC. See [“Upgrading the CMC”](#) (page 26).
- Upgrade storage systems in one or more management groups or available storage systems. See [“Upgrading storage systems in a management group or available storage systems”](#) (page 26).

Setting upgrade preferences

Upgrade preferences, opened from the Help menu, control the settings described in [Table 4](#) (page 25).

Table 4 Upgrade preferences

Preference	Options
Download	<ul style="list-style-type: none">• Automatic—System automatically downloads upgrade files when they become available.• On Request—(Default) When upgrades are available, you must click Start Download from the Upgrades tab to download upgrade files.
Bandwidth Speed	<ul style="list-style-type: none">• Fast—Uses available bandwidth for file downloads.• Normal—(Default) Reduces bandwidth used for file downloads (downloads are about 20% slower).
Upgrade Selection Mode	<ul style="list-style-type: none">• Advanced—Lets you select which upgrade or patch to install.• Normal—(Default) Installs all upgrades and patches needed to upgrade to the latest version.
Download Directory	Use the default directory or click Browse to select a different one. If you change the directory, the system copies all files from the old directory to the new one and downloads any additional files you need.

Checking for available upgrades

A button on the CMC menu bar lets you see when upgrades are available. The button changes to the following states:

- **Software Up to Date**—System has checked for upgrades, and no upgrades are available.
- **Upgrades Available**—Upgrades are available. Click for more information.
- **Check for Upgrades**—Appears briefly until the system connects to the FTP site to check for upgrades. Persists when the system cannot access the FTP site to check for upgrades. Click to see a message on the Upgrades tab.

To check for upgrades:

1. Click **Upgrades Available** or **Check for Upgrades**.

The Upgrades tab opens, with the list of available upgrades. If you are using VSS, the HP DSM for MPIO, or the CLI, it also lists the most current versions of those programs.

The system accesses an FTP site and determines which patch and upgrade files you need. If you do not have all of the files, the **Start Download** button is displayed.

If you see a message about a problem connecting to the download site, check your Internet connection and click **Check for Upgrades**.

If the system you are working from does not have Internet access, use a different system to download the files and copy them to a location the first system can access. Then from the CMC Upgrade tab, click **Use Local Media** to see the available upgrades.

2. If the button is displayed, click **Start Download**.

Patch and upgrade files download to your download directory. To verify or change the download directory, see [“Setting upgrade preferences” \(page 25\)](#).

If you decide not to install some upgrades now, the next time you check for upgrades, the system will not download files you already have.

3. Continue with the upgrading process.

To upgrade the CMC, see [“Upgrading the CMC” \(page 26\)](#).

To upgrade storage systems in a management group or available systems, see [“Upgrading storage systems in a management group or available storage systems” \(page 26\)](#).

Upgrading the CMC

When you upgrade the CMC, the CMC closes, then installs the upgrade.

Before upgrading the CMC, verify the following:

- You are using version 9.0 or later of the CMC.
- You are using the CMC that you want to upgrade.

To upgrade the CMC:

1. When you see an available CMC upgrade, click **Install** next to the CMC icon.
The CMC closes, then installs the upgrade.
2. Follow the instructions in the installation wizard.
If you want to install the SNMP MIBs, select the Complete installation.

Upgrading storage systems in a management group or available storage systems

The system upgrades one storage system at a time in management group, making every effort to keep volumes available. The process includes all patches and upgrades.

When upgrading storage systems in the Available Systems list, all storage systems in the list are upgraded.

Before upgrading the storage systems in a management group or available storage systems, verify the following:

- For all storage systems:
 - You are using version 9.0 or later of the CMC.
 - All storage systems are in a good state, with no critical events.

- If you are using the HP DSM for MPIO or the CLI, upgrade them first.
- If you are using VSS, be prepared to upgrade it during the SAN/iQ upgrade process.
- For storage systems in a management group:
 - You are logged in to the management groups you want to upgrade.
 - At least two storage systems are running managers, so you will not lose quorum.
 - If you want to maintain availability to volumes during the upgrade, your management group must be configured with Network RAID-10 (2-Way Mirror) or higher.
 - All virtual managers in the management group are stopped.

NOTE: During the upgrade procedure, you may receive a warning that the CPU Utilization value exceeds 90, for example: CPU Utilization = 97.8843. Value exceeds 90. This is an expected occurrence during an upgrade. No action is needed.





To upgrade storage systems:

1. When you see an upgrade for a management group or for available systems, click **Install** next to the item you want to upgrade.
2. If the Package Selection window opens, select the upgrade or patch you want to install, and click **OK**.
 The Package Selection window opens if you have your Upgrade Select Mode set to Advanced. See [“Setting upgrade preferences” \(page 25\)](#).
 The Install Action Required window opens, with a list of software that must be upgraded before continuing with the SAN/iQ upgrade. The current versions were downloaded to your download directory.
3. If you are using the listed software, verify that the current versions are installed, or install the current versions
4. After verifying or installing the listed upgrades, select the check box at the bottom of the window, and click **Continue**.
 To stop the SAN/iQ upgrade process while you verify or install the listed software, click **Cancel**.
 The Upgrade Progress window opens, showing the upgrade progress for each storage system, then the Upgrade Summary window opens. To see information about the upgrades, click **Installation Succeeded**. To save the upgrade information to a file, click **Export Information**.
 The Upgrade Summary also window lists software that must be upgraded before finishing the SAN/iQ upgrade. For any software listed, the current version was downloaded to your download directory.
5. If you are using the listed software, verify that the current versions are installed or install the current versions.
6. After verifying or installing the listed upgrades, select the check box at the bottom of the window, and click **Finish**.

Monitoring upgrade progress

During the upgrade, the status of the installation and upgrades is tracked on the Upgrade Progress window. The Upgrade Progress window displays the storage systems that are being upgraded and the activity on each system.

Table 5 Status icons for upgrade progress

Icon	Description
	Copying files to system
	Installing on system
	System is waiting
	System upgrade is complete

Reviewing the upgrade summary

The Upgrade Summary window displays the results of the upgrade installations.

- If the installation failed, click **Installation Failed** to view more information about the failure.
- To save the upgrade summary information, click **Export Information**.

Registering advanced features for a storage system

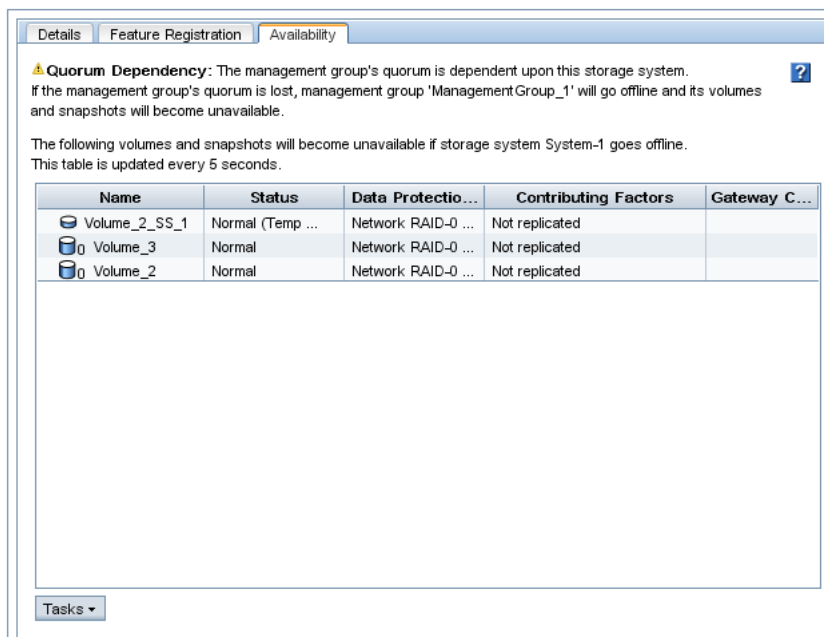
Using the Feature Registration tab, register individual storage systems for advanced features.

For more information about registering advanced features, see [“Registering advanced features” \(page 222\)](#).

Determining volume and snapshot availability

The Availability tab displays which volumes' and snapshots' availability depends on this storage system staying online. Details include the data protection level and what factors contribute to the availability status, such as the status of storage systems participating in any replication or a RAID restripe in progress.

Figure 7 Availability tab



Checking status of dedicated boot devices

Some storage systems contain either one or two dedicated boot devices. Dedicated boot devices may be compact flash cards or hard drives. If a storage system has dedicated boot devices, the Boot Devices tab appears in the Storage configuration category. Storage systems that do not have dedicated boot devices will not display the Boot Devices tab.

In storage systems with two dedicated boot devices, both devices are active by default. If necessary, compact flash cards can be deactivated or activated using the buttons on this tab. However, you should only take action on these cards if instructed by HP Technical Support.

Checking boot device status

View dedicated boot device status in the Boot Devices tab window in the Storage category in the storage system tree.

Getting there

1. Select a storage system in the navigation window and log in if necessary.
2. Open the tree below the storage system and select **Storage**.
3. Select the **Boot Devices** tab.

The status of each dedicated boot device on the storage system is listed in the Status column. [Table 6 \(page 29\)](#) describes the possible status for boot devices.

NOTE: Some statuses only occur in a storage system with two boot devices.

Table 6 Boot device status

Boot device status	Description
Active	The device is synchronized and ready to be used.
Inactive	The device is ready to be removed from the storage system. It will not be used to boot the storage system.
Failed	The device encountered an I/O error and is not ready to be used.
Unformatted	The device has not yet been used in a storage system. It is ready to be activated.
Not Recognized	The device is not recognized as a boot device.
Unsupported	The device cannot be used. (For example, the compact flash card is the wrong size or type.)

NOTE: When the status of a boot device changes, an event is generated. See [“Alarms and events overview” \(page 85\)](#).

Replacing a dedicated boot device

If a boot hard drive fails, you will see an event that the boot device is faulty. Replace it with a new drive. The boot device drives support hot swapping and do not require activation.

3 Storage Configuration: Disk RAID and Disk Management

Use the Storage configuration category to configure and manage RAID and individual disks for storage systems.

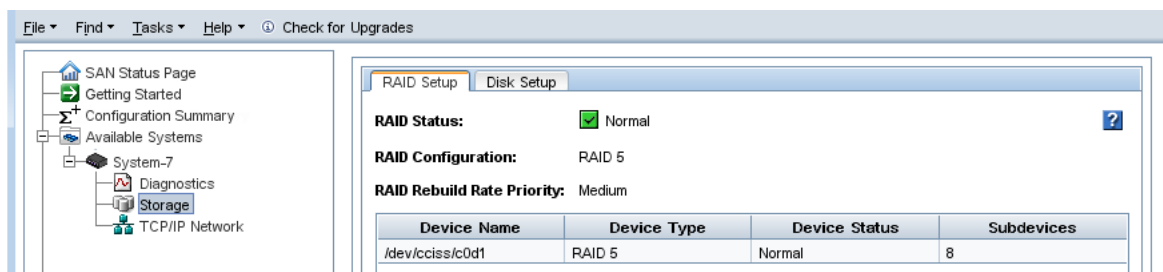
Configuring RAID and managing disks

For each storage system, you can select the RAID configuration and the RAID rebuild options, and monitor the RAID status. You can also review disk information and, for some models, manage individual disks.

Getting there

1. In the navigation window, select a storage system and log in if necessary.
2. Open the tree under the storage system and select the **Storage** category.

Figure 8 Viewing the storage configuration category for a storage system



Columns in the RAID Setup tab show four categories:

- Device Name
- Device Type or the RAID level
- Device Status
- Subdevices

Status indicators

On the RAID Setup tab and the Disk Setup tab, the text or icon color indicates status. [Table 7 \(page 30\)](#) lists the status and color indicators for three categories.

- RAID Device Status
- Disk Status
- Disk Health

Table 7 Status and color definitions

Status	Color
Normal	Green
Inactive	Yellow / orange
Uninitialized	Yellow
Rebuilding	Blue
Off or Removed	Red
Marginal	Yellow
Faulty	Red

Table 7 Status and color definitions *(continued)*

Status	Color
Hot Spare	Green
Hot Spare Down	Yellow

Configuring and managing RAID

Managing the RAID settings of a storage system includes:

- Choosing the right RAID configuration for your storage needs
- Setting or changing the RAID configuration, if necessary
- Setting the rate for rebuilding RAID
- Monitoring the RAID status for the storage system
- Reconfiguring RAID when necessary

RAID Levels

The availability of certain RAID levels is determined by the number of storage system hard drives.

Table 8 Descriptions of RAID levels

RAID level	Description
RAID 0 – Striping (No Fault Tolerance)	Offers the greatest capacity and performance without data protection. If you select this option, you will experience data loss if a hard drive that holds the data fails. However, because no logical drive capacity is used for redundant data, this method offers the best capacity. This method offers the best processing speed by reading two stripes on different hard drives at the same time and by not having a parity drive.
RAID 1 – Mirroring	Offers a good combination of data protection and performance. RAID 1 or drive mirroring creates fault tolerance by storing duplicate sets of data on a minimum of two hard drives. There must be an even number of drives for RAID 1. RAID 1 and RAID 1+0(10) are the most costly fault tolerance methods because they require 50 percent of the drive capacity to store the redundant data. RAID 1 mirrors the contents of one hard drive in the array onto another. If either hard drive fails, the other hard drive provides a backup copy of the files and normal system operations are not interrupted.
RAID 1+0 – Mirroring and Striping	Offers the best combination of data protection and performance. RAID 1+0 or drive mirroring creates fault tolerance by storing duplicate sets of data on a minimum of four hard drives. There must be an even number of drives for RAID 1+0. RAID 1+0(10) and RAID 1 are the most costly fault tolerance methods because they require 50 percent of the drive capacity to store the redundant data. RAID 1+0(10) first mirrors each drive in the array to another, and then stripes the data across the mirrored pair. If a physical drive fails, the mirror drive provides a backup copy of the files and normal system operations are not interrupted. RAID 1+0(10) can withstand multiple simultaneous drive failures, as long as the failed drives are not mirrored to each other.
RAID 5	Offers the best combination of data protection and usable capacity while also improving performance over RAID 6. RAID 5 stores parity data across all the physical drives in

Table 8 Descriptions of RAID levels *(continued)*

RAID level	Description
	the array and allows more simultaneous read operations and higher performance. If a drive fails, the controller uses the parity data and the data on the remaining drives to reconstruct data from the failed drive. The system continues operating with a slightly reduced performance until you replace the failed drive. RAID 5 can only withstand the loss of one drive without total array failure. It requires an array with a minimum of three physical drives. Usable capacity is $N-1$ where N is the number of physical drives in the logical array.
RAID 6	Offers the best data protection and is an extension of RAID 5. RAID 6 uses multiple parity sets to store data and can therefore tolerate up to 2 drive failures simultaneously. RAID 6 requires a minimum of 4 drives. Performance is lower than RAID 5 due to parity data updating on multiple drives. RAID 6 uses two disk for parity; its fault tolerance allows two disks to fail simultaneously. Usable capacity is $N-2$ where N is the number of physical drives in the logical array.

Explaining RAID devices in the RAID setup report

In the Storage category, the RAID Setup tab lists the RAID devices in the storage system and provides information about them. An example of the RAID setup report is shown in [Figure 9 \(page 32\)](#). Information listed in the report is described in [Table 9 \(page 32\)](#).

Figure 9 RAID setup report

Device Name	Device Type	Device Status	Subdevices
/dev/cciss/c0d1	RAID 5	Normal	8

RAID devices by RAID type

Each RAID type creates different sets of RAID devices. [Table 9 \(page 32\)](#) contains a description of the variety of RAID devices created by the different RAID types as implemented on various storage systems.

Table 9 Information in the RAID setup report

This item	Describes this
Device Name	The disk sets used in RAID. The number and names of devices varies by storage system and RAID level.
Device Type	The RAID level of the device. For example, in a P4300 G2, RAID 5 displays a Device Type of RAID 5 and subdevices as 8.
Device Status	The RAID status of the device.
Subdevices	The number of disks included in the device.

Virtual RAID devices

If you are using the VSA, the only RAID available is virtual RAID. After installing the VSA, virtual RAID is configured automatically if you first configured the data disk in the VI Client.

HP recommends installing VMware ESX Server on top of a server with a RAID 5 or RAID 6 configuration.

Planning the RAID configuration

The RAID configuration you choose for the storage system depends on your plans for data fault tolerance, data availability, and capacity growth.

-
- ⚠ CAUTION:** Plan your RAID configuration carefully. After you have configured RAID, you cannot change the RAID configuration without deleting all data on the storage system.
-

Data protection

Keeping multiple copies of your data ensures that data is safe and remains available in the case of disk failure. There are two ways to achieve data protection:

- Configure RAID 1, RAID 10, RAID 5, RAID 5 + spare, RAID 50, or RAID 6 within each storage system to ensure data redundancy.
- Always use Network RAID to mirror data volumes across storage systems in a cluster, regardless of RAID level, for added data protection and high availability.

Using RAID for data redundancy

Within each storage system, RAID 1 or RAID 10 ensures that two copies of all data exist. If one of the disks in a RAID pair goes down, data reads and writes continue on the other disk. Similarly, RAID 5, RAID 50, or RAID 6 provides redundancy by spreading parity evenly across the disks in the set.

If one disk in a RAID 5 set, or two disks in a RAID 6 set goes down, data reads and writes continue on the remaining disks in the set. In RAID 50, up to one disk in each RAID 5 set can go down, and data reads and writes continue on the remaining disks.

RAID protects against failure of disks within a storage system, but not against failure of an entire storage system. For example, if network connectivity to the storage system is lost, then data reads and writes to the storage system cannot continue.

NOTE: If you plan on using clusters with only a single storage system, use RAID 1 and RAID 10, RAID 5, or RAID 6 to ensure data redundancy within that storage system.

Using Network RAID in a cluster

A cluster is a group of storage systems across which data can be protected by using Network RAID. Network RAID protects against the failure of a RAID disk set within a storage system, failure of an entire storage system or external failures like networking or power. For example, if an entire storage system in a cluster becomes unavailable, data reads and writes continue because the missing data can be obtained from the other storage systems.

Using disk RAID with Network RAID in a cluster

Always use Network RAID in a cluster to protect volumes across storage systems. The redundancy provided by RAID 10, RAID 5, RAID 50, or RAID 6 ensures availability at the storage system level. Using Network RAID for volumes in a cluster ensures availability at the cluster level. For example:

- Using Network RAID, up to three copies of a volume can be created on a cluster of three storage systems. The Network RAID configuration ensures that two of the three storage systems can go offline and the volume is still accessible.
- Configuring RAID 10 on these storage systems means that each of these three copies of the volume is stored on two disks within the storage system, for a total of six copies of each volume. For a 50 GB volume, 300 GB of disk capacity is used.

RAID 5 and RAID 50 use less disk capacity than RAID 1 or RAID 10, so they can be combined with Network RAID and still use capacity efficiently. One benefit of configuring RAID 5 or RAID 50 in storage systems that use Network RAID in a cluster is that if a single disk goes down, the data on that storage system can be rebuilt using RAID instead of requiring a complete copy from another storage system in the cluster. Rebuilding the disks within a single set is faster and creates less of a performance impact to applications accessing data than copying data from another storage system in the cluster.

RAID 6 provides similar space benefits to RAID 5, with the additional protection of being able to survive the loss of up to two drives.

NOTE: If you are protecting volumes across a cluster, configuring the storage system for RAID 1 or RAID 10 consumes half the capacity of the storage system. Configuring the storage system for RAID 5 or RAID 50 provides redundancy within each storage system while allowing most of the disk capacity to be used for data storage. RAID 6 provides greater redundancy on a single storage system, but consumes more disk space than RAID 5.

Table 10 (page 34) summarizes the differences in data availability and safety of the different RAID levels on stand-alone storage systems compared with those RAID levels with Network RAID configured volumes in a cluster.

Table 10 Data availability and safety in RAID configurations

Configuration	Data safety and availability during disk failure	Data availability if entire storage system fails or if network connection to storage system lost
Stand-alone storage systems, RAID 0	No	No
Stand-alone storage systems, RAID 1, RAID 10, RAID 10 + spare	Yes. In any configuration, 1 disk per mirrored pair can fail.	No
Stand-alone storage systems, RAID 5, RAID 5 + spare, RAID 50	Yes, for 1 disk per array	No
Stand-alone storage systems, RAID 6	Yes, for 2 disks per array	No
Volumes configured with Network RAID-10 or greater on clustered storage systems, RAID 0	Yes. However, if any disk in the storage system fails, the entire storage system must be copied from another storage system in the cluster.	Yes
Volumes configured with Network RAID-10 or greater on clustered storage systems, RAID 5, RAID 50	Yes. 1 disk per RAID set can fail without copying from another storage system in the cluster.	Yes

Table 10 Data availability and safety in RAID configurations *(continued)*

Configuration	Data safety and availability during disk failure	Data availability if entire storage system fails or if network connection to storage system lost
Volumes configured with Network RAID-10 or greater on clustered storage systems, RAID 6	Yes. 2 disks per RAID set can fail without copying from another storage system in the cluster.	Yes
Volumes configured with Network RAID-10 or greater on clustered VSAs with virtual RAID	Depends on the underlying RAID configuration of the storage system on which the VSA is installed. HP recommends configuring RAID 5 or RAID 6.	Yes, if underlying storage system configured for RAID other than RAID 0.

Mixing RAID configurations

You may mix storage systems with different configurations of RAID within a cluster. This allows you to add new storage systems with different RAID levels. However, be certain to calculate the capacity of additional storage systems configured with the desired RAID level, because the cluster operates at the smallest usable per-storage system capacity.

For instance, your SAN uses four 12 TB HP LeftHand P4500s configured with RAID 10. You purchase two additional 12 TB HP LeftHand P4500s which you want to configure with RAID 5.

In the existing cluster, a single 12 TB HP LeftHand P4500 configured with RAID 10 provides 6 TB of usable storage. A single 12 TB HP LeftHand P4500 configured with RAID 5 provides 9 TB of usable storage. However, due to the restrictions of how the cluster uses capacity, the 12 TB HP LeftHand P4500 configured with RAID 5 will be limited to 6 TB per storage system.

In general, the best practice is to avoid mixing configurations of various numbers or capacities of drives, so that the SAN fully utilizes the available capacity of each cluster.

Setting RAID rebuild rate

Choose the rate at which the RAID configuration rebuilds if a disk is replaced.

NOTE: The RAID rebuild rate cannot be set on a VSA, since there is no physical hardware to rebuild.

General guidelines for setting the RAID rebuild rate

Use the following guidelines when deciding where to set the RAID rebuild rate.

- Setting the rate high is preferred for rebuilding RAID quickly and protecting data. However, it slows down user access to data.
- Setting the rate low allows users quicker access to data during the rebuild, but slows the rebuild rate.

Setting the RAID rebuild rate

1. In the navigation window, log in to a storage system and select the Storage category.
2. On the RAID Setup tab, click **RAID Setup Tasks** and select the **RAID Rebuild Rate Priority** choice.
3. Change the rebuild settings as desired on the RAID Rebuild Rate Priority window.
4. Click **OK**.

The settings are then ready when a RAID rebuild takes place.

Reconfiguring RAID

Reconfiguring RAID on a storage system or a VSA destroys any data stored on that storage system. For VSAs, there is no alternate RAID choice, so the only outcome for reconfiguring RAID is to wipe out all data.

- Changing preconfigured RAID on a new storage system
RAID must be configured on individual storage systems before they are added to a management group. To change the preconfigured RAID level of a storage system, make the change before you add the storage system to a management group.
- Changing RAID on storage systems in management groups
You cannot reconfigure RAID on a storage system that is already in a management group. To change the RAID configuration for a storage system that is in a management group, you must first remove it from the management group.

⚠ CAUTION: Changing the RAID configuration will erase all the data on the disks.

To reconfigure RAID

1. In the navigation window, log in to the storage system and select the **Storage** category.
2. On the RAID Setup tab, click **RAID Setup Tasks** and select **Reconfigure RAID**.
3. Select the RAID configuration from the list.
4. Click **OK**.
5. Click **OK** on the message that opens.
RAID starts configuring.

NOTE: A storage system may take several hours for the disks to synchronize in a RAID 10, RAID 5, RAID 50, or RAID 6 configuration. During this time, performance will be degraded. When the RAID status on the RAID Setup tab shows Normal, the disks provide fully operational data redundancy, and performance returns to normal.

Monitoring RAID status

RAID is critical to the operation of the storage system. If RAID has not been configured, the storage system cannot be used. Monitor the RAID status of a storage system to ensure that it remains normal. If the RAID status changes, a CMC event is generated. For more information about events and event notification, see [“Alarms and events overview” \(page 85\)](#).

Data reads and writes and RAID status

A RAID status of Normal, Rebuild, or Degraded all allow data reads and writes. The only time data cannot be written to and read from the storage system is if the RAID status shows Off.

Data redundancy and RAID status

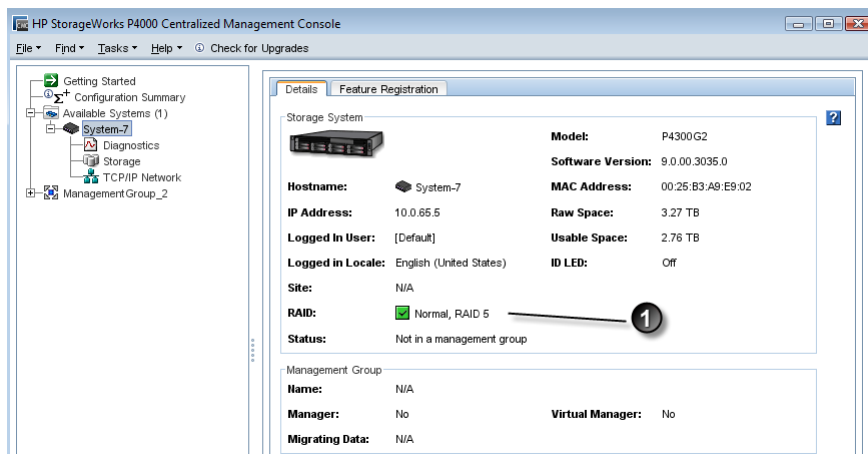
In a RAID 1, RAID 10, RAID 5 or RAID 50 configuration, when RAID is degraded, there is no full data redundancy. Therefore, data is at additional risk if there is a disk failure when RAID is degraded.

In RAID 6, when RAID is degraded due to a single drive failure, the data is still not at risk for a second failure. However, if it is degraded due to the failure of two drives, then data would be at risk if another drive failed.

- △ **CAUTION:** In a degraded RAID 1 or RAID 10 configuration, loss of a second disk within a pair results in data loss. In a degraded RAID 5 configuration, loss of a second disk results in data loss. In a degraded RAID 50 configuration, loss of a second disk in a single RAID 5 set results in data loss. In a degraded RAID 6 configuration, the loss of three drives results in data loss.

The RAID status is located at the top of the RAID Setup tab in Storage. RAID status also appears in the Details tab on the main CMC window when a storage system is selected in the navigation window.

Figure 10 Monitoring RAID status on the main CMC window



1. RAID status

The status displays one of four RAID states.

- **Normal**—RAID is synchronized and running. No action is required.
- **Rebuilding**—A new disk has been inserted in a drive bay, or a hot spare has been activated, and RAID is currently rebuilding. No action is required.
- **Degraded**—RAID is degraded. A disk may have failed or have been removed from its bay. For hot-swap storage systems, simply replace the faulty, inactive, uninitialized, or missing disk.
- **Off**—Data cannot be stored on the storage system. The storage system is offline and flashes in the navigation window.
- **None**—RAID is unconfigured.

Managing disks

Use the Disk Setup tab to monitor disk information and perform disk management tasks as listed in Table 11 (page 37).

- △ **CAUTION:** Hot-swapping drives is not supported for RAID 0 on any storage system.

Table 11 Disk management tasks for storage systems

Disk setup function	Model where available
Monitor disk information	All

Getting there

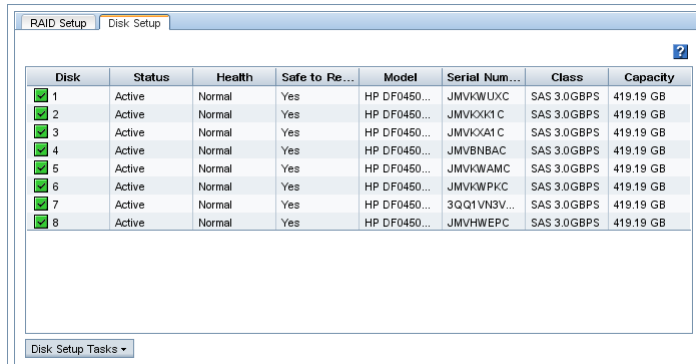
1. In the navigation window, select a storage system.
2. Select the **Storage** category in the tree below it.
3. Select the **Disk Setup** tab.

Reading the disk report on the Disk Setup tab

The Disk Setup tab provides a status report of the individual disks in a storage system.

Figure 11 (page 38) shows the Disk Setup tab and Table 12 (page 39) describes the corresponding disk report.

Figure 11 Example of columns in the Disk Setup tab



The screenshot shows a software interface with two tabs: "RAID Setup" and "Disk Setup". The "Disk Setup" tab is active, displaying a table with 8 columns: Disk, Status, Health, Safe to Re..., Model, Serial Num..., Class, and Capacity. There are 8 rows of data, each representing a disk. Each row starts with a green checkmark in the "Disk" column. The "Status" column contains "Active" for all disks. The "Health" column contains "Normal" for all disks. The "Safe to Re..." column contains "Yes" for all disks. The "Model" column contains "HP DF0450..." for all disks. The "Serial Num..." column contains various serial numbers. The "Class" column contains "SAS 3.0GBPS" for all disks. The "Capacity" column contains "419.19 GB" for all disks. Below the table is a "Disk Setup Tasks" button.

Disk	Status	Health	Safe to Re...	Model	Serial Num...	Class	Capacity
1	Active	Normal	Yes	HP DF0450...	JMVKWUXC	SAS 3.0GBPS	419.19 GB
2	Active	Normal	Yes	HP DF0450...	JMVKX1C	SAS 3.0GBPS	419.19 GB
3	Active	Normal	Yes	HP DF0450...	JMVKX1C	SAS 3.0GBPS	419.19 GB
4	Active	Normal	Yes	HP DF0450...	JMVBNBAC	SAS 3.0GBPS	419.19 GB
5	Active	Normal	Yes	HP DF0450...	JMVKWAMC	SAS 3.0GBPS	419.19 GB
6	Active	Normal	Yes	HP DF0450...	JMVKWPKC	SAS 3.0GBPS	419.19 GB
7	Active	Normal	Yes	HP DF0450...	3QQ1VN3V...	SAS 3.0GBPS	419.19 GB
8	Active	Normal	Yes	HP DF0450...	JMVHWEPC	SAS 3.0GBPS	419.19 GB

Table 12 Description of items on the disk report

Column	Description
Disk	Corresponds to the physical slot in the storage system.
Status	Whether the disk is <ul style="list-style-type: none">• Active (on and participating in RAID)• Uninitialized (is not part of an array)• Inactive (is part of an array, and on, but not participating in RAID)• Off or removed• Hot spare (for RAID configurations that support hot spares)
Health	Drive health is one of the following <ul style="list-style-type: none">• Normal• Marginal (predictive failure status indicating “replace as soon as convenient”)• Faulty (predictive failure status indicating “replace immediately”)
Safe to Remove	Indicates if it is safe to hot-remove a disk.
Model	The model of the disk.
Serial Number	The serial number of the disk.
Class	The class (type) of disk, for example, SATA 3.0 GB.
Capacity	The data storage capacity of the disk.

Verifying disk status

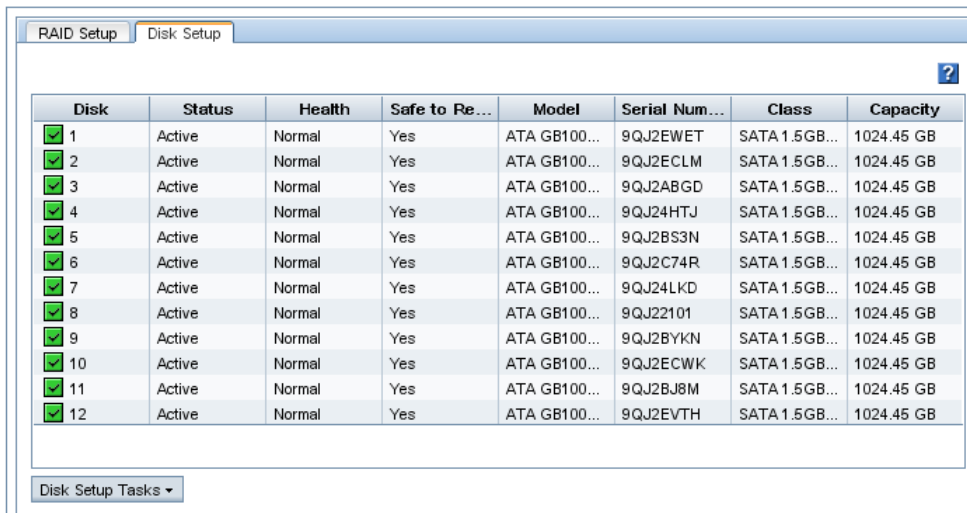
Check the Disk Setup window to determine the status of disks and to take appropriate action on individual disks when preparing to replace them.

Viewing disk status for the DL320s (NSM 2120)

The disks are labeled 1 through 12 in the Disk Setup window, shown in [Figure 12 \(page 40\)](#), and correspond to the disk drives from left to right (1-4-7-10 on the top row, and 2-5-8-11 on the second row and so on), as shown in [Figure 13 \(page 40\)](#) when looking at the front of the DL320s (NSM 2120).

For the DL320s (NSM 2120), the columns Health and Safe to Remove help assess the health of a disk and tell you whether you can replace it without losing data.

Figure 12 Viewing the Disk Setup tab in a DL320s (NSM 2120)



Disk	Status	Health	Safe to Re...	Model	Serial Num...	Class	Capacity
1	Active	Normal	Yes	ATA GB100...	9QJ2EWET	SATA 1.5GB...	1024.45 GB
2	Active	Normal	Yes	ATA GB100...	9QJ2ECLM	SATA 1.5GB...	1024.45 GB
3	Active	Normal	Yes	ATA GB100...	9QJ2ABGD	SATA 1.5GB...	1024.45 GB
4	Active	Normal	Yes	ATA GB100...	9QJ24HTJ	SATA 1.5GB...	1024.45 GB
5	Active	Normal	Yes	ATA GB100...	9QJ2BS3N	SATA 1.5GB...	1024.45 GB
6	Active	Normal	Yes	ATA GB100...	9QJ2C74R	SATA 1.5GB...	1024.45 GB
7	Active	Normal	Yes	ATA GB100...	9QJ24LKD	SATA 1.5GB...	1024.45 GB
8	Active	Normal	Yes	ATA GB100...	9QJ22101	SATA 1.5GB...	1024.45 GB
9	Active	Normal	Yes	ATA GB100...	9QJ2BYKN	SATA 1.5GB...	1024.45 GB
10	Active	Normal	Yes	ATA GB100...	9QJ2ECWK	SATA 1.5GB...	1024.45 GB
11	Active	Normal	Yes	ATA GB100...	9QJ2BJ8M	SATA 1.5GB...	1024.45 GB
12	Active	Normal	Yes	ATA GB100...	9QJ2EVTH	SATA 1.5GB...	1024.45 GB

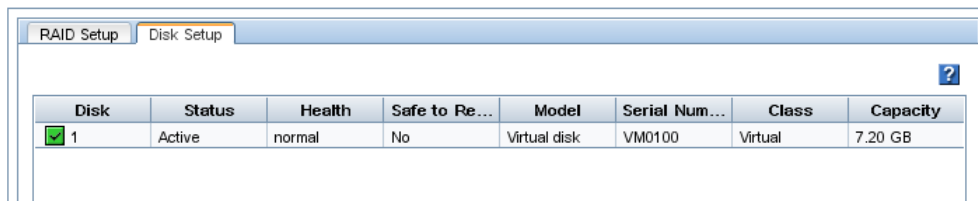
Figure 13 Diagram of the drive bays in a DL320s (NSM 2120)



Viewing disk status for the VSA

For the VSA, the Disk Setup window shows 1 virtual disk.

Figure 14 Viewing the disk status of a VSA



Disk	Status	Health	Safe to Re...	Model	Serial Num...	Class	Capacity
1	Active	normal	No	Virtual disk	VM0100	Virtual	7.20 GB

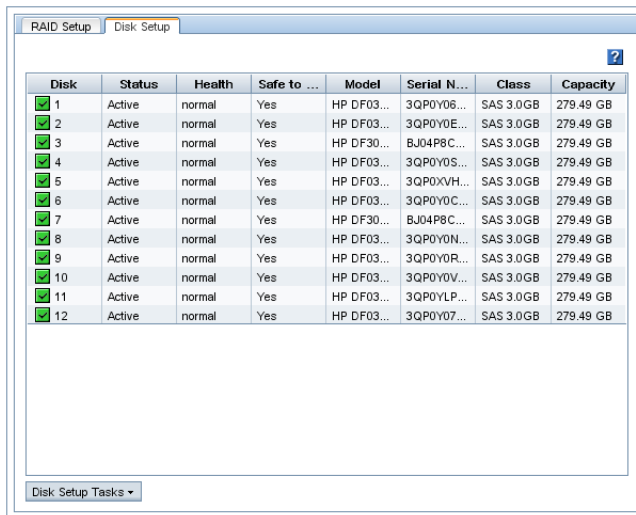
NOTE: To change the size of the data disk in a VSA, see the *HP P4000 VSA Installation and Configuration Guide* for instructions about recreating the disk in the VI Client.

Viewing disk status for the HP LeftHand P4500 and HP P4500 G2

The disks are labeled 1 through 12 in the Disk Setup window, shown in [Figure 15 \(page 41\)](#), and correspond to the disk drives from left to right (1-4-7-10 on the top row, and 2-5-8-11 on the second row and so on), as shown in [Figure 16 \(page 41\)](#), when you are looking at the front of the HP LeftHand P4500 and HP P4500 G2.

For the HP LeftHand P4500 and HP P4500 G2, the columns Health and Safe to Remove help you assess the health of a disk and tell you whether or not you can replace it without losing data.

Figure 15 Viewing the Disk Setup tab in a HP LeftHand P4500 and HP P4500 G2



Disk	Status	Health	Safe to ...	Model	Serial N...	Class	Capacity
1	Active	normal	Yes	HP DF03...	3QP0Y06...	SAS 3.0GB	279.49 GB
2	Active	normal	Yes	HP DF03...	3QP0Y0E...	SAS 3.0GB	279.49 GB
3	Active	normal	Yes	HP DF30...	BJ04P8C...	SAS 3.0GB	279.49 GB
4	Active	normal	Yes	HP DF03...	3QP0Y0S...	SAS 3.0GB	279.49 GB
5	Active	normal	Yes	HP DF03...	3QP0XVH...	SAS 3.0GB	279.49 GB
6	Active	normal	Yes	HP DF03...	3QP0Y0C...	SAS 3.0GB	279.49 GB
7	Active	normal	Yes	HP DF30...	BJ04P8C...	SAS 3.0GB	279.49 GB
8	Active	normal	Yes	HP DF03...	3QP0Y0N...	SAS 3.0GB	279.49 GB
9	Active	normal	Yes	HP DF03...	3QP0Y0R...	SAS 3.0GB	279.49 GB
10	Active	normal	Yes	HP DF03...	3QP0Y0V...	SAS 3.0GB	279.49 GB
11	Active	normal	Yes	HP DF03...	3QP0YLP...	SAS 3.0GB	279.49 GB
12	Active	normal	Yes	HP DF03...	3QP0Y07...	SAS 3.0GB	279.49 GB

Figure 16 Diagram of the drive bays in a HP LeftHand P4500 and HP P4500 G2

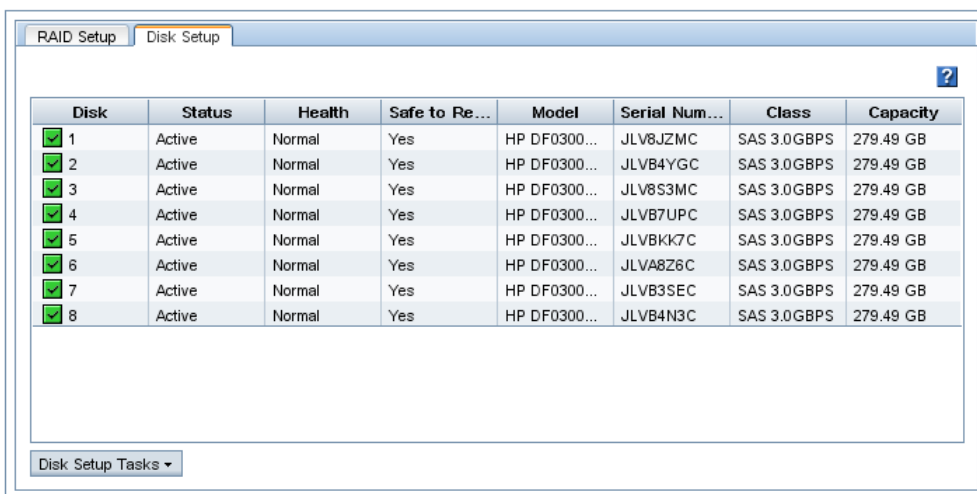


Viewing disk status for the HP LeftHand P4300 and HP P4300 G2

The disks are labeled 1 through 8 in the Disk Setup window, shown in [Figure 17 \(page 41\)](#), and correspond to the disk drives from left to right (1, 3, 5, and 7 on the top row, and 2, 4, 6, and 8 on the second row), as shown in [Figure 18 \(page 41\)](#), when you are looking at the front of the HP LeftHand P4300 and HP P4300 G2.

For the P4300 and the P4300 G2, the columns Health and Safe to Remove help you assess the health of a disk and tell you whether or not you can replace it without losing data.

Figure 17 Viewing the Disk Setup tab in a HP LeftHand P4300 and HP P4300 G2



Disk	Status	Health	Safe to Re...	Model	Serial Num...	Class	Capacity
1	Active	Normal	Yes	HP DF0300...	JLV8JZMC	SAS 3.0GBPS	279.49 GB
2	Active	Normal	Yes	HP DF0300...	JLVB4YGC	SAS 3.0GBPS	279.49 GB
3	Active	Normal	Yes	HP DF0300...	JLV8S3MC	SAS 3.0GBPS	279.49 GB
4	Active	Normal	Yes	HP DF0300...	JLVB7UPC	SAS 3.0GBPS	279.49 GB
5	Active	Normal	Yes	HP DF0300...	JLVBKK7C	SAS 3.0GBPS	279.49 GB
6	Active	Normal	Yes	HP DF0300...	JLVA8Z6C	SAS 3.0GBPS	279.49 GB
7	Active	Normal	Yes	HP DF0300...	JLVB3SEC	SAS 3.0GBPS	279.49 GB
8	Active	Normal	Yes	HP DF0300...	JLVB4N3C	SAS 3.0GBPS	279.49 GB

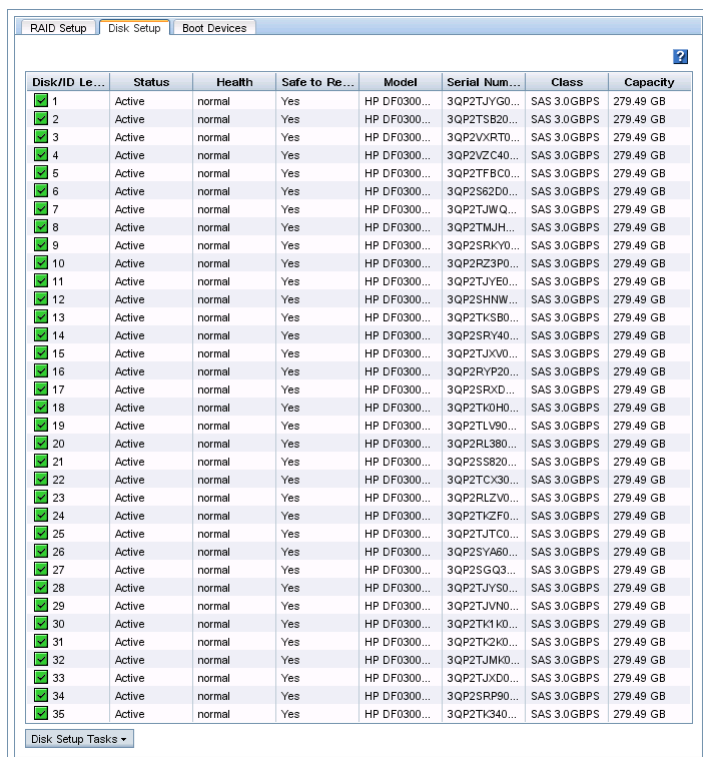
Figure 18 Diagram of the drive bays in a HP LeftHand P4300 and HP P4300 G2



Viewing disk status for the HP LeftHand P4800 and HP P4800 G2


The disks are labeled 1 through 35 in the Disk Setup window, shown in [Figure 19 \(page 42\)](#), and correspond to the disk drives from left to right (1-8-15-22 on the top row, and 2-9-16-23-30 on the second row and so on), as shown in [Figure 20 \(page 42\)](#), when you are looking at the front of the HP LeftHand P4800 and HP P4800 G2.

Figure 19 Viewing the Disk Setup tab in a HP LeftHand P4800 and HP P4800 G2



Disk/ID Le...	Status	Health	Safe to Re...	Model	Serial Num...	Class	Capacity
1	Active	normal	Yes	HP DF0300...	3QP2TJYG0...	SAS 3.0GBPS	279.49 GB
2	Active	normal	Yes	HP DF0300...	3QP2TSB20...	SAS 3.0GBPS	279.49 GB
3	Active	normal	Yes	HP DF0300...	3QP2VXRT0...	SAS 3.0GBPS	279.49 GB
4	Active	normal	Yes	HP DF0300...	3QP2VZC40...	SAS 3.0GBPS	279.49 GB
5	Active	normal	Yes	HP DF0300...	3QP2TFBC0...	SAS 3.0GBPS	279.49 GB
6	Active	normal	Yes	HP DF0300...	3QP2S6ZD0...	SAS 3.0GBPS	279.49 GB
7	Active	normal	Yes	HP DF0300...	3QP2TJWQ...	SAS 3.0GBPS	279.49 GB
8	Active	normal	Yes	HP DF0300...	3QP2TMJH...	SAS 3.0GBPS	279.49 GB
9	Active	normal	Yes	HP DF0300...	3QP2SRKY0...	SAS 3.0GBPS	279.49 GB
10	Active	normal	Yes	HP DF0300...	3QP2RZ3P0...	SAS 3.0GBPS	279.49 GB
11	Active	normal	Yes	HP DF0300...	3QP2TJYE0...	SAS 3.0GBPS	279.49 GB
12	Active	normal	Yes	HP DF0300...	3QP2SHNW...	SAS 3.0GBPS	279.49 GB
13	Active	normal	Yes	HP DF0300...	3QP2TKS80...	SAS 3.0GBPS	279.49 GB
14	Active	normal	Yes	HP DF0300...	3QP2SRV40...	SAS 3.0GBPS	279.49 GB
15	Active	normal	Yes	HP DF0300...	3QP2TJXV0...	SAS 3.0GBPS	279.49 GB
16	Active	normal	Yes	HP DF0300...	3QP2RYP20...	SAS 3.0GBPS	279.49 GB
17	Active	normal	Yes	HP DF0300...	3QP2SRXD...	SAS 3.0GBPS	279.49 GB
18	Active	normal	Yes	HP DF0300...	3QP2TK0H0...	SAS 3.0GBPS	279.49 GB
19	Active	normal	Yes	HP DF0300...	3QP2TLV90...	SAS 3.0GBPS	279.49 GB
20	Active	normal	Yes	HP DF0300...	3QP2RL380...	SAS 3.0GBPS	279.49 GB
21	Active	normal	Yes	HP DF0300...	3QP2SS820...	SAS 3.0GBPS	279.49 GB
22	Active	normal	Yes	HP DF0300...	3QP2TCX30...	SAS 3.0GBPS	279.49 GB
23	Active	normal	Yes	HP DF0300...	3QP2RLZV0...	SAS 3.0GBPS	279.49 GB
24	Active	normal	Yes	HP DF0300...	3QP2TKZF0...	SAS 3.0GBPS	279.49 GB
25	Active	normal	Yes	HP DF0300...	3QP2TJTC0...	SAS 3.0GBPS	279.49 GB
26	Active	normal	Yes	HP DF0300...	3QP2SYA60...	SAS 3.0GBPS	279.49 GB
27	Active	normal	Yes	HP DF0300...	3QP2SGQ3...	SAS 3.0GBPS	279.49 GB
28	Active	normal	Yes	HP DF0300...	3QP2TJY50...	SAS 3.0GBPS	279.49 GB
29	Active	normal	Yes	HP DF0300...	3QP2TJVM0...	SAS 3.0GBPS	279.49 GB
30	Active	normal	Yes	HP DF0300...	3QP2TK1K0...	SAS 3.0GBPS	279.49 GB
31	Active	normal	Yes	HP DF0300...	3QP2TK2K0...	SAS 3.0GBPS	279.49 GB
32	Active	normal	Yes	HP DF0300...	3QP2TJMK0...	SAS 3.0GBPS	279.49 GB
33	Active	normal	Yes	HP DF0300...	3QP2TJXD0...	SAS 3.0GBPS	279.49 GB
34	Active	normal	Yes	HP DF0300...	3QP2SRP90...	SAS 3.0GBPS	279.49 GB
35	Active	normal	Yes	HP DF0300...	3QP2TK340...	SAS 3.0GBPS	279.49 GB

Figure 20 Diagram of the drive bays in a HP LeftHand P4800 and HP P4800 G2



Replacing a disk

The correct procedure for replacing a disk in a storage system depends upon a number of factors, including the RAID configuration, the data protection level of volumes and snapshots, and the number of disks being replaced. Replacing a disk in a storage system that is in a cluster requires rebuilding data either just on the replaced disk or, in the case of RAID 0, on the entire storage system.

Replacing a disk in a storage system includes the following basic steps.

- Planning for rebuilding data on either the disk or the entire storage system (all storage systems)
- Powering the disk off in the CMC (non-hot-swap storage systems)
- Physically replacing the disk in the storage system (all storage systems)
- Powering the disk on in the CMC (non-hot-swap storage systems)
- Rebuilding RAID on the disk or on the storage system (all storage systems)

Table 13 (page 43) lists disk replacement requirements for specific configurations and storage systems.

Table 13 Disk replacement requirements

Storage system or configuration	Requirements
Hot-swap storage systems configured for RAID 1, 10, 5, 50, or 6	RAID is normal and Safe to Remove status is yes. See “Replacing disks in hot-swap storage systems” (page 43) .
VSA	Replace disk on host server according to manufacturer's instructions.
RAID 0 configuration	Plan for data conservation before replacing disk. Power off disk in CMC before physically replacing disk. See “Replacing a disk in RAID 0” (page 45) .

Additional information about preparing for disk replacement is included in the following sections below:

- [“Preparing for a disk replacement” \(page 44\)](#)
- [“Best practice checklist for single disk replacement in RAID 0” \(page 44\)](#)
- [“Best practice checklist for single disk replacement in RAID 1, RAID 10, RAID 5, RAID 50, and RAID 6 ” \(page 45\)](#)

Using Repair Storage System

Repair Storage System is a procedure that allows you to replace a disk and trigger only one resync of data, rather than a complete restripe. Repair Storage System creates a place-holder system in the cluster, while allowing the storage system needing repair to be removed for the disk replacement. See [“Repairing a storage system” \(page 138\)](#) for more information.

In the following circumstances, you may have to use the Repair Storage System feature when replacing disks.

- When RAID is OFF on a storage system with RAID 0
- When replacing multiple disks on a storage system with RAID 5, RAID 50 or RAID 6
- When multiple disks on the same mirror set need to be replaced on a storage system with RAID 10.

Replacing disks in hot-swap storage systems

In hot-swap storage systems configured with RAID 1, RAID 10, RAID 5, RAID 50, or RAID 6, a faulty or failed disk can be removed and replaced with a new one. RAID will rebuild and the drive will return to Normal status.

⚠ CAUTION: Before replacing a drive in a hot-swap storage system, always check the Safe to Remove status to verify that the drive can be removed without causing RAID to go Off.

When RAID is Normal in RAID 1, RAID 10, RAID 5, RAID 50, or RAID 6, all drives indicate they are safe to remove. However, you should only replace one drive at a time. If it is necessary to replace more than one drive, always check the Safe to Remove status again. Wait up to two

minutes for the status to fully update before checking it again. If the status indicates the second drive is safe to remove, then it can be replaced.

For example, if an array is Rebuilding, no other drives in the array (except for unused hot-spare drives) are safe to remove. However, if the configuration includes two or more arrays and those arrays are Normal, the Safe To Remove status indicates that drives in those other arrays may be replaced.

NOTE: The Safe To Remove status is always No when in a RAID 0 configuration until the drive is powered off. Hot Spare, Inactive, and Uninitialized drives are always safe to remove.

Preparing for a disk replacement

Use this section to replace a single disk under the following conditions:

- You know which disk needs to be replaced through SAN/iQ monitoring.
- When viewed in the Disk Setup tab, the Drive Health column shows Marginal (replace as soon as possible) or Faulty (replace right away).
- RAID is still on, though it may be degraded and a drive is inactive.

Use the instructions in [“Replacing disks appendix” \(page 241\)](#) for these situations:

- If RAID is off
- If you are unsure which disk to replace

The instructions in the appendix include contacting Customer Support for assistance in either identifying the disk that needs to be replaced or, for replacing more than one disk, the sequence in which they should be replaced.

To prepare for disk replacement

Preparing for a disk replacement differs according to the RAID level of the storage system and whether it is a hot-swap storage system. Carefully plan any disk replacement to ensure data safety, regardless of whether the storage system is hot-swap. The following checklists outline steps to help ensure data remains safe while replacing a disk.

Identify physical location of storage system and disk

Before beginning the disk replacement process, identify the physical location of both the storage system in the rack and the disk in the storage system.

- Know the name and physical location of the storage system that needs the disk replacement.
- Know the physical position of the disk in the storage system. See [“Verifying disk status” \(page 39\)](#) for diagrams of disk layout in the various storage systems.
- Have the replacement disk ready and confirm that it is of the right size and has the right carrier.

Best practice checklist for single disk replacement in RAID 0



CAUTION: Do not use hot-swap procedures on any storage system configured with in RAID 0. Removing a drive in a RAID 0 configuration results in data loss.

In RAID 0, always power off the drive in the CMC before removing it. RAID 0 provides no fault tolerance by itself, so when you power off the drive, the data on the storage system is lost. Therefore, if you need to replace a disk in a RAID 0 configuration, HP recommends the following:

- All volumes and snapshots have a Network RAID level that provides redundancy across storage systems. Any level other than Network RAID-0 provides this kind of redundancy.
- If volumes or snapshots are not protected (they are configured for Network RAID 0), change them to a Network RAID level that provides data protection. You will have to wait for the data to restripe, which could take some time.
- If the cluster does not have enough space for changing the Network RAID level, take a backup of the volumes or snapshots and then delete them from the cluster.
After the disk replacement is complete, you can recreate the volumes and restore the data from the backup.
- All volumes and snapshots show a status of Normal.
- Any volumes or snapshots that are being deleted have finished deleting.
- Use the instructions in [“Replacing disks appendix” \(page 241\)](#) to replace more than one disk, or if you are unsure which disk needs replacing.

Best practice checklist for single disk replacement in RAID 1, RAID 10, RAID 5, RAID 50, and RAID 6

There are no prerequisites for this case; however, HP recommends that:

- All volumes and snapshots show a status of Normal.
- Any volumes or snapshots that were being deleted have completed deletion.
- One of the following:
 - RAID status is Normal
 - If RAID is Rebuilding or Degraded, for storage systems that support hot-swapping of drives, the Safe to Remove column indicates Yes (the drive can safely be replaced).

Replacing a disk in RAID 0

Complete the following checklist for single disk replacement RAID 0.

Manually power off the disk in the CMC for RAID 0

First power off the disk you are replacing in the CMC, which causes RAID to go off.

1. In the navigation window, select the storage system containing the disk to be replaced.
2. Select the **Storage** category.
3. Select the **Disk Setup** tab.
4. Select the disk in the list to power off.
5. Click **Disk Setup Tasks**, and select **Power Off Disk**.
6. Click **OK** on the confirmation message.

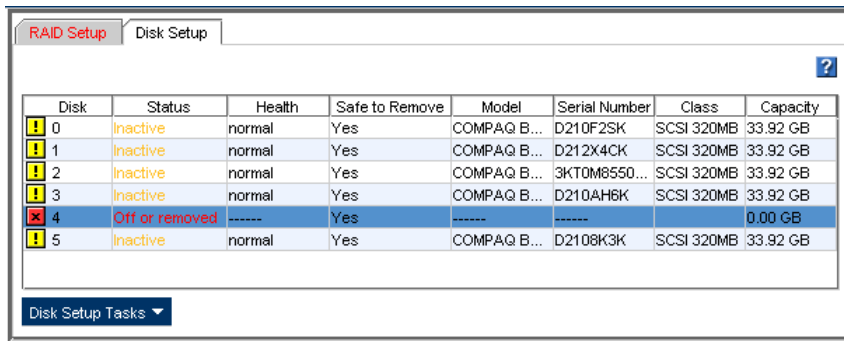
Physically replace the disk drive in the storage system

See the hardware documentation for the storage system.

Manually powering on the disk in the CMC

After inserting the new disk into the storage system, the disk must be powered on from the Storage category Disk Setup tab. Until the disk is powered on, it is listed as Off or Missing in the Status column and the other columns display dotted lines. [Figure 21 \(page 46\)](#) shows a representation of a missing disk in a storage system.

Figure 21 Viewing a power off or missing disk



Disk	Status	Health	Safe to Remove	Model	Serial Number	Class	Capacity
0	Inactive	normal	Yes	COMPAQ B...	D210F2SK	SCSI 320MB	33.92 GB
1	Inactive	normal	Yes	COMPAQ B...	D212X4CK	SCSI 320MB	33.92 GB
2	Inactive	normal	Yes	COMPAQ B...	3KT0M8550...	SCSI 320MB	33.92 GB
3	Inactive	normal	Yes	COMPAQ B...	D210AH6K	SCSI 320MB	33.92 GB
4	Off or removed	-----	Yes	-----	-----	-----	0.00 GB
5	Inactive	normal	Yes	COMPAQ B...	D2108K3K	SCSI 320MB	33.92 GB

Disk Setup Tasks ▼

1. In the navigation window, select the storage system with the replaced drive.
2. Select the **Storage** category in the navigation tree.
3. Click the **Disk Setup** tab.
4. Select the disk in the list to power on.
5. Click **Disk Setup Tasks**, and select **Power On Disk**.
6. Click **OK** on the confirmation message.

Volume restriping

After the disk is powered on, RAID changes to Normal. Volumes start restriping on the entire storage system. Note that there may be a delay of up to a couple of minutes before you can see that volumes are restriping.

Replacing a disk in a hot-swap storage system

The hot-swap storage systems are listed below.

- DL320s [NSM 2120]
- HP LeftHand P4300 and P4500
- HP P4500 G2 and P4500 G2
- HP P4800

Complete the checklist for replacing a disk in RAID 1, RAID 10, RAID 5, RAID 50, or RAID 6. Then follow the appropriate procedures for the storage system.

Replace the disk

You may remove and replace a disk from these hot-swap storage systems after checking that the Safe to Remove status indicates “Yes” for the drive to be replaced.

Physically replace the disk drive in the storage system

See the hardware documentation that came with your storage system for information about physically replacing disk drives in the storage system.

RAID rebuilding

After the disk is replaced, RAID starts rebuilding on the replaced disk. There may be a delay of up to a couple of minutes before you can see that RAID is rebuilding on the RAID Setup or Disk Setup tabs.

Figure 22 RAID rebuilding on the RAID Setup tab

RAID Setup
Disk Setup
Boot Devices

RAID Status: Rebuilding

RAID Configuration: RAID 5 (6 disks)

RAID Rebuild Rate Percent: 100%

Device Name	Device Type	Device Status	Subdevices
/dev/scsi/host0/bus1/target0/lun...	RAID 5	Rebuilding: 7% complete, estimating 82	6
/dev/scsi/host1/lun1/target0/lun...	RAID 5	Normal	6

RAID Setup Tasks

Figure 23 Disk rebuilding on the Disk Setup tab

RAID Setup
Disk Setup
Boot Devices

Disk	Status	Health	Safe to Remove	Model	Serial Number	Class	Capacity
1	Active	normal	No	ST3250823	3ND1CXGG	SATA 3.0GB	232.74 GB
2	Active	normal	No	ST3250823	4ND0JQHN	SATA 3.0GB	232.74 GB
3	Rebuilding	normal	Yes	ST3250823	3ND1DDDF	SATA 3.0GB	232.74 GB
4	Active	normal	No	ST3250823	4ND0JRF8	SATA 3.0GB	232.74 GB
5	Active	normal	No	ST3250823	3ND1DGP6	SATA 3.0GB	232.74 GB
6	Active	normal	No	ST3250823	4ND0LVYS	SATA 3.0GB	232.74 GB
7	Active	normal	Yes	ST3250823	4ND0M6H2	SATA 3.0GB	231.90 GB
8	Active	normal	Yes	ST3250823	4ND0LNXL	SATA 3.0GB	231.90 GB
9	Active	normal	Yes	ST3250823	4ND0MEDX	SATA 3.0GB	231.90 GB
10	Active	normal	Yes	ST3250823	4ND0LNY1	SATA 3.0GB	231.90 GB
11	Active	normal	Yes	ST3250823	4ND0LZXA	SATA 3.0GB	231.90 GB
12	Active	normal	Yes	ST3250823	4ND0LP2Q	SATA 3.0GB	231.90 GB

Disk Setup Tasks

4 Managing the network

A physical storage system has two TCP/IP network interfaces (NICs). For each physical storage system you can:

- Configure the individual TCP/IP interfaces, including settings for speed and duplex, frame size, and NIC flow control.
- Bond NICs to ensure continuous network access or to improve bandwidth.
- Set up and manage DNS servers, and a network routing table.
- Manage the SAN/iQ communication interface, and update the list of managers running in the management group to which a storage system belongs.

NOTE: The VSA has only one network interface and does not support changing the following items:

- NIC bonding
 - NIC flow control
 - Frame size
 - TCP interface speed or duplex
-

Network best practices

- Isolate the SAN, including CMC traffic, on a separate network. If the SAN must run on a public network, use a VPN to secure data and CMC traffic.
- Configure all the network characteristics on a storage system before creating a management group, or before adding the storage system to a management group and cluster.
- Use static IP addresses, or reserved addresses if using DHCP.
- Configure storage system settings for speed and duplex, frame size, and flow control BEFORE bonding NICs and before putting the storage system into a management group and cluster.
- If adding a second IP address to a storage system, the second IP address must be on a separate subnet. If the two IP addresses are on the same subnet, they must be bonded.

Split network configurations

Implement a split network configuration to use the CMC on a management network separate from the storage or iSCSI network. HP recommends that you only configure a management network when the storage systems are equipped with a sufficient number of network adapters so that the storage systems can support both bonded NICs for optimal redundancy and performance on the storage network, along with a separate management interface. Keep in mind the following requirements when considering the creation of a management network.

- Configure a management network only if your storage systems have two or more NICs. Two NICs are the minimum number required to create a bonded interface for the storage network to provide maximum redundancy and performance.
- Use only available P4000 upgrade kits to alter the hardware configuration of a P4000 storage system. Alterations using any other hardware are not supported.
- Be aware that the CMC can manage storage systems via a routed network or by installation of a second NIC on the storage network on the CMC host. This is the most common method for CMC management.

- When configuring a management interface on a P4000 storage system, make sure that interface is on a separate network. Configuring two separate interfaces on the same IP network is not supported and will result in communication problems in the cluster.
- When configuring a management interface on a P4000 storage system, be aware that only one interface can be configured with a default gateway, and that interface should be the management interface. Configuring two default gateways will result in communication problems in the cluster.
- When configuring a management interface on a P4000 storage system, you must designate the storage interface as the SAN/iQ interface for that storage system in the CMC. This is done on the Communications tab in the TCP/IP configuration category for that storage system.

Changing network configurations

Changing the network configuration of a storage system may affect connectivity with the network and application servers. Consequently, we recommend that you configure network characteristics on individual storage systems before creating a management group or adding them to existing clusters.

If you do need to change the network characteristics of a storage system while it is in a cluster, be sure to follow HP's recommendations. See [“Best practices when changing network characteristics”](#) (page 49).

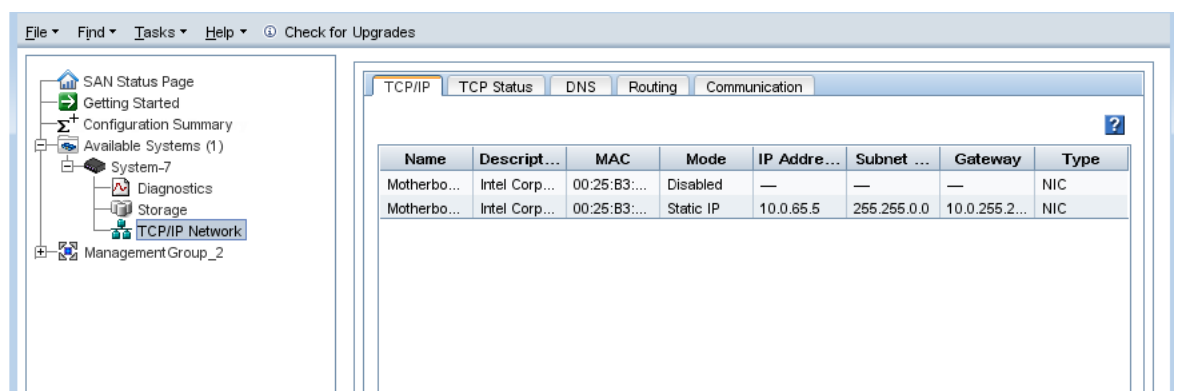
Best practices when changing network characteristics

- Plan to make network changes during off-peak hours to minimize the impact of those changes.
- Make changes on one storage system at a time.
- Some network changes cause the storage server to restart the SAN/iQ services, making the storage system unavailable for a short time. Check the Availability tab for each storage system to see if any volumes will become unavailable if the services restart on the storage system.
Volumes and snapshots may become temporarily unavailable while services restart. Examples include unreplicated volumes, or snapshots that are causing a restripe of the data.
- After the changes are in place, verify the iSCSI sessions. You may need to update the sessions.

Getting there

1. In the navigation window, select a storage system.
2. Open the tree under the storage system, and select **TCP/IP Network**.

Figure 24 Viewing the TCP/IP Network category for a storage system



Managing settings on network interfaces

Configure or change the settings of the network interfaces in the storage systems. See [“Network best practices”](#) (page 48) for more information.

Requirements

These settings must be configured before creating NIC bonds.

TCP status tab

Review the status of the TCP interfaces. Change the speed and duplex, frame size, and NIC flow control of an interface. These changes can only take place on interfaces that are not in a bond.

NOTE: You cannot change the speed, duplex, frame size, or flow control of a VSA.

Review the status of the network interfaces on the TCP Status tab.

Table 14 Status of and information about network interfaces

Column	Description
Name	Name of the interface. Entries vary depending on the storage system. <ul style="list-style-type: none">• bond0—The bonded interface(s) (appears only if storage system is configured for bonding)• Motherboard:Port1• Motherboard:Port2• G4-Motherboard:Port1• G4-Motherboard:Port2• Eth0
Description	Describes each interface listed. For example, the bond0 is the Logical Failover Device.
Speed/Method	Lists the actual operating speed reported by the device.
Duplex/Method	Lists duplex as reported by the device.
Status	Describes the state of the interface. See “NIC status in Active-Passive configuration” (page 57) for a detailed description of individual NIC status.
Frame Size	Lists the frame size setting for the device.
Preferred	(For Active-Passive bonds) Indicates whether the device is set as preferred. The preferred interface is the interface within an Active-Passive bond that is used for data transfer during normal operation.

Changing speed and duplex settings

The settings for the storage system and the switch must be the same. Available settings are listed in [Table 15](#) (page 50).

Table 15 Setting storage system speed and duplex settings

Storage system setting speed/duplex	Switch setting speed/duplex
Auto/Auto	Auto/Auto
1000/Full	1000/Full
100/Full	100/Full

Table 15 Setting storage system speed and duplex settings *(continued)*

Storage system setting speed/duplex	Switch setting speed/duplex
100/Half	100/Half
10/Full	10/Full
10/Half	10/Half

NOTE: The VSA does not support changing the speed and duplex settings.

Requirements

- These settings must be configured before creating NIC bonds.
- If you change these settings, you must ensure that both sides of the NIC cable are configured in the same manner. For example, if the storage system is set for Auto/Auto, the switch must be set the same.
- If you edit the speed or duplex on a disabled or failed NIC, the new setting will not be applied until the NIC is enabled or connectivity is restored.

Best practice

Change the speed and duplex settings while the storage system is in the Available Systems pool and not in a management group.

To change the speed and duplex

1. In the navigation window, select the storage system and log in.
2. Open the tree, and select **TCP/IP Network**.
3. Select the **TCP Status** tab in the tab window.
4. Select the interface to edit.
5. Click **TCP/IP Status Tasks**, and select **Edit**.
6. Select the combination of speed and duplex that you want.
7. Click **OK**.

A series of status messages appears. Then the changed setting appears in the TCP status report.

NOTE: You can also use the Configuration Interface to edit the TCP speed and duplex. See “Setting the TCP speed, duplex, and frame size” (page 237).

Changing NIC frame size

Network frame size affects management data traffic and replication. Use the same network frame size in all storage systems in a cluster to ensure consistent data traffic and replication.

Requirements

If you plan to change the frame size, that change must be configured before creating NIC bonds.

Best practices

Change the frame size while the storage system is in the Available Systems pool and not in a management group.

The frame size specifies the size of data packets that are transferred over the network. The default Ethernet standard frame size is 1500 bytes. The maximum allowed frame size is 9000 bytes.

Increasing the frame size improves data transfer speed by allowing larger packets to be transferred over the network and by decreasing the CPU processing time required to transfer data. However,

increasing the frame size requires that routers, switches, and other devices on your network support that frame size.

NOTE: Increasing the frame size can cause decreased performance and other network problems if routers, switches, or other devices on your network do not support frame sizes greater than 1500 bytes. If you are unsure about whether your routers and other devices support larger frame sizes, keep the frame size at the default setting.

If you edit the frame size on a disabled or failed NIC, the new setting will not be applied until the NIC is enabled or connectivity is restored.

To avoid potential connectivity and performance problems with other devices on your network, keep the frame size at the default setting. The frame size on the storage system should correspond to the frame size on Windows and Linux application servers. If you decide to change the frame size, set the same frame size on all storage systems on the network, and set compatible frame sizes on all clients that access the storage systems.

Consult with your network administrator for recommended storage system frame sizes and the corresponding frame sizes in bytes for Windows and Linux clients in your environment.

Jumbo frames

Frame sizes that are greater than 1500 bytes are called jumbo frames. Jumbo frames must be supported and configured on each Windows or Linux client accessing the storage system and also on each network switch between the storage system and the Windows or Linux clients.

Jumbo frames can co-exist with 1500 byte frames on the same subnet if the following conditions are met:

- Every device downstream of the storage system on the subnet must support jumbo frames.
- If you are using 802.1q virtual LANs, jumbo frames and nonjumbo frames must be segregated into separate VLANs.

NOTE: The frame size for a bonded logical interface must be equal to the frame size of the NICs in the bond.

Editing the NIC frame size

To edit the frame size:

1. In the navigation window, select a storage system and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **TCP Status** tab.
4. Select the interface to edit.
5. Click **TCP Status Tasks**, and select **Edit**.
6. Select **Set To** in the Frame Size section.
7. Enter a value between 1500 and 9000 bytes in the Set To field.
8. Click **OK**.

A series of status messages display. Then the changed setting appears in the TCP status report.

NOTE: You can also use the Configuration Interface to edit the frame size.

Changing NIC flow control

Set the network flow control settings the same on every storage system in the cluster. Enable flow control on the NICs to prevent data transmission overruns that result in packets being dropped. With flow-control enabled, network packets that would otherwise be dropped will not have to be retransmitted.

NOTE: The VSA does not support changing flow control settings.

Requirements

- These settings must be configured before creating NIC bonds.
- All NICs should have (or must have, if they are bonded) the same flow control settings.
- Flow control cannot be changed when the port is disabled.

Enabling NIC flow control

To enable NIC flow control:

1. In the navigation window, select a storage system, and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **TCP Status** tab.
4. Select the interface to edit.
5. Click **TCP Status Tasks**, and select **Edit**.
6. Select **On** to enable the flow control on the NIC.
7. Click **OK**.
8. Repeat steps 4 through 7 for all the NICs you want to enable.

When you have enabled flow control on both NICs and then you bond those NICs, the NIC flow control column shows the physical NICs as enabled and the bond0 as disabled. However, flow control is enabled and working in this case.

The TCP/IP tab

Lists the network interfaces on the storage system. On the TCP/IP tab, you can bond interfaces, disable an interface, configure an IP address, and ping servers from the storage system.


The TCP/IP tab

Identifying the network interfaces

A storage system comes with two Ethernet interfaces. To use either interface, you must connect an Ethernet cable to either port and configure the interface in the Configuration Interface or the CMC. These ports are named and labeled on the back of the storage system.

[Table 16 \(page 54\)](#) lists the methods to identify the NICs. You can work with the NICs in the CMC or through the Configuration Interface which is accessed through the storage system's serial port, as described in ["Using the Configuration Interface" \(page 235\)](#).

Table 16 Identifying the network interfaces on the storage system

Ethernet interfaces	
Where labeled	Label
In the TCP/IP Network configuration category in the CMC <ul style="list-style-type: none"> • TCP/IP tab • TCP Status tab 	<ul style="list-style-type: none"> • eth0, eth1 • Motherboard:Port0, Motherboard:Port1 • G4-Motherboard:Port1, G4-Motherboard:Port2 • Motherboard:Port1, Motherboard:Port2 For bonded interfaces: <ul style="list-style-type: none"> • BondN or Bond0
In the Configuration Interface available through the storage system's serial port	<ul style="list-style-type: none"> • Intel Gigabit Ethernet • Broadcom Gigabit Ethernet
On the label on the back of the storage system	<ul style="list-style-type: none"> • Eth0, Eth1 • Represented by a graphical symbol similar to the symbols below:
	

Pinging an IP address

Because the SAN should be on a private network, you can ping target IP addresses from a storage system using the CMC. You can ping from any enabled interface listed on the TCP/IP tab. You can ping any IP address, such as an iSCSI server or an SNMP monitor server.

To ping an IP address

1. Select a storage system, and open the tree below it.
2. Select the **TCP/IP Network** category.
3. Select the **TCP/IP Tasks** menu, and select **Ping** from the menu.
4. Select which Network Interface to ping from, if you have more than one enabled.
A bonded interface has only one interface from which to ping.
5. Enter the IP address to ping, and click **Ping**.
If the server is available, the ping is returned in the Ping Results window.
If the server is not available, the ping fails in the Ping Results window.

Configuring the IP address manually

Use the TCP/IP Network category to add or change the IP address for a network interface.

1. Select a storage system, and open the tree below it.
2. Select **TCP/IP Network** category, and click the **TCP/IP** tab.
3. Select the interface from the list for which to configure or change the IP address.
4. Click **Edit**.
5. Select IP address, and complete the fields for IP address, Subnet mask, and Default gateway.
6. Click **OK**.
7. Click **OK** on the confirmation message.
8. Click **OK** on the message notifying you of the automatic log out.

NOTE: Wait a few moments for the IP address change to take effect.

9. Log in to the newly addressed storage system.

Using DHCP

A DHCP server becomes a single point of failure in your system configuration. If the DHCP server goes offline, then IP addresses may be lost.

- △ **CAUTION:** If you use DHCP, be sure to reserve statically assigned IP addresses for all storage systems on the DHCP server. This is required, because management groups use unicast communication.

NOTE: If you plan to bond the NICs, you must use a static IP address.

To set IP address using DHCP

1. Select from the list the interface to configure for use with DHCP.
2. Click **Edit**.
3. Select **Obtain an address automatically using the DHCP/BOOTP protocol**.
4. Click **OK**.
5. Click **OK** on the confirmation message.
6. Click **OK** on the message notifying you of the automatic log out.

NOTE: Wait a few moments for the IP address change to take effect.

Configuring network interface bonds

To ensure consistent failover characteristics and traffic distribution, use the same network bond type in all the storage systems in a cluster. Network interface bonding provides high availability, fault tolerance, load balancing and/or bandwidth aggregation for the network interface cards in the storage system. Bonds are created by joining physical NICs into a single logical interface. This logical interface acts as the master interface, controlling and monitoring the physical slave interfaces.

Bonding two interfaces for failover provides fault tolerance at the local hardware level for network communication. Failures of NICs, Ethernet cables, individual switch ports, and/or entire switches can be tolerated while maintaining data availability. Bonding two interfaces for aggregation provides bandwidth aggregation and localized fault tolerance. Bonding the interfaces for load balancing provides both load balancing and localized fault tolerance.

NOTE: The VSA does not support NIC bonding.

Depending on your storage system hardware, network infrastructure design, and Ethernet switch capabilities, you can bond NICs in one of three ways:

- **Active-Passive.** You specify a preferred NIC for the bonded logical interface to use. If the preferred NIC fails, then the logical interface begins using another NIC in the bond until the preferred NIC resumes operation. When the preferred NIC resumes operation, data transfer resumes on the preferred NIC.
- **Link Aggregation Dynamic Mode.** The logical interface uses both NICs simultaneously for data transfer. This configuration increases network bandwidth, and if one NIC fails, the other

continues operating normally. To use Link Aggregation Dynamic Mode, your switch must support 802.3ad.

△ **CAUTION:** Link Aggregation Dynamic Mode requires plugging both NICs into the same switch. This bonding method does not protect against switch failure.

- **Adaptive Load Balancing (ALB).** The logical interface balances data transmissions through both NICs to enhance the functionality of the server and the network. Adaptive Load Balancing automatically incorporates fault tolerance features as well.

Best practices

- Adaptive Load Balancing is the recommended bonding method, as it combines the benefits of the increased transmission rates of 802.3ad with the network redundancy of Active-Passive. Adaptive Load Balancing does not require additional switch configurations.
- Verify and, if necessary, change the Speed, Duplex, Frame Size, and Flow Control settings for both interfaces that you plan to bond.
- Link Aggregation Dynamic Mode does not protect against switch failure, because both NICs must be plugged into the same switch. Link Aggregation Dynamic Mode provides bandwidth gains, because data is transferred over both NICs simultaneously. For Link Aggregation Dynamic Mode, both NICs must be plugged into the same switch, and that switch must be LACP-capable, and both support and be configured for 802.3ad aggregation.
- For Active-Passive, plug the two NICs on the storage system into separate switches. While Link Aggregation Dynamic Mode will only survive a port failure, Active-Passive will survive a switch failure.

IP address for NIC bonds

Allocate a static IP address for the logical bond interface (bond0). You cannot use DHCP for the bond IP.

NIC bonding and speed, duplex, frame size, and flow control settings

These settings are controlled on the TCP/IP and TCP Status tabs of the TCP/IP Network configuration category. If you change these settings, you must ensure that *both* sides of the NIC cable are configured in the same manner. For example, if the storage system is set for Auto/Auto, the switch must be set the same. See “TCP status tab” (page 50) for more information.

Table 17 Comparison of Active-Passive, link aggregation dynamic mode, and Adaptive Load Balancing bonding

Feature	Active-Passive	Link aggregation dynamic mode	Adaptive load balancing
Bandwidth	Use of 1 NIC at a time provides normal bandwidth.	Simultaneous use of both NICs increases bandwidth.	Simultaneous use of both NICs increases bandwidth.
Protection during port failure	Yes	Yes	Yes
Protection during switch failure	Yes. NICs can be plugged into different switches.	No. Both NICs are plugged into the same switch.	Yes. NICs can be plugged into different switches.
Requires support for 802.3ad link aggregation?	No	Yes	No

How Active-Passive bonding works

Bonding NICs for Active-Passive allows you to specify a preferred interface that will be used for data transfer. This is the active interface. The other interface acts as a backup, and its status is "Passive (Ready)."

Physical and logical interfaces

The two NICs in the storage system are labeled as listed in [Table 18 \(page 57\)](#). If both interfaces are bonded for failover, the logical interface is labeled bond0 and acts as the master interface. As the master interface, bond0 controls and monitors the two slave interfaces which are the physical interfaces.

Table 18 Bonded network interfaces

Failover name	Failover description
bond0	Logical interface acting as master
eth0 or Motherboard:Port1	Physical interface acting as slave
eth1 or Motherboard:Port2	Physical interface acting as slave

The logical master interface monitors each physical slave interface to determine if its link to the device to which it is connected, such as a router, switch, or repeater, is up. As long as the interface link remains up, the interface status is preserved.

Table 19 NIC status in Active-Passive configuration

If the NIC status is	The NIC is
Active	Currently enabled and in use
Passive (Ready)	Slave to a bond and available for failover
Passive (Failed)	Slave to a bond and no longer has a link

If the active NIC fails, or if its link is broken due to a cable failure or a failure in a local device to which the NIC cable is connected, then the status of the NIC becomes Passive (Failed) and the other NIC in the bond, if it has a status of Passive (Ready), becomes active.

This configuration remains until the failed preferred interface is brought back online. When the failed interface is brought back online, it becomes Active. The other NIC returns to the Passive (Ready) state.

Requirements for Active-Passive

To configure Active-Passive:

- Both NICs should be enabled.
- NICs should be connected to separate switches.

Which physical interface is preferred

When the Active-Passive bond is created, if both NICs are plugged in, the SAN/iQ software interface becomes the active interface. The other interface is Passive (Ready).

For example, if Eth0 is the preferred interface, it will be active and Eth1 will be Passive (Ready). Then, if Eth0 fails, Eth1 changes from Passive (Ready) to active. Eth0 changes to Passive (Failed).

Once the link is fixed and Eth0 is operational, there is a 30-second delay and then Eth0 becomes the active interface. Eth1 returns to the Passive (Ready) state.

NOTE: When the active interface comes back up, there is a 30-second delay before it becomes active.

Table 20 Example Active-Passive failover scenario and corresponding NIC status

Example failover scenario	NIC status
1. Active-Passive bond0 is created. The active (preferred) interface is Eth0.	<ul style="list-style-type: none"> Bond0 is the master logical interface. Eth0 is Active. Eth1 is connected and is Passive (Ready).
2. Active interface fails. Bond0 detects the failure and Eth1 takes over.	<ul style="list-style-type: none"> Eth0 status becomes Passive (Failed). Eth1 status changes to Active.
3. The Eth0 link is restored.	<ul style="list-style-type: none"> Eth0 status changes to Active after a 30 second delay. Eth1 status changes to Passive (Ready).

Summary of NIC status during failover

Table 21 (page 58) shows the states of Eth0 and Eth1 when configured for Active-Passive.


Table 21 NIC status during failover with Active-Passive

Failover status	Status of Eth0	Status of Eth1
Normal Operation	Preferred: Yes Status: Active Data Transfer: Yes	Preferred: No Status: Passive (Ready) Data Transfer: No
Eth0 Fails, Data Transfer Fails Over to Eth1	Preferred: Yes Status: Passive (Failed) Data Transfer: No	Preferred: No Status: Active Data Transfer: Yes
Eth0 Restored	Preferred: Yes Status: Active Data Transfer: Yes	Preferred: No Status: Passive (Ready) Data Transfer: No

Example network cabling topologies with Active-Passive

Two simple network cabling topologies using Active-Passive in high availability environments are shown in Figure 25 (page 59).


Figure 25 Active-Passive in a two-switch topology with server failover



1. Servers
2. HP P4000
3. Storage cluster
4. GigE trunk
5. Active path
6. Passive path

The two-switch scenario in [Figure 25 \(page 59\)](#) is a basic, yet effective, method for ensuring high availability. If either switch fails, or a cable or NIC on one of the storage systems fails, the Active-Passive bond causes the secondary connection to become active and take over.

Figure 26 Active-Passive failover in a four-switch topology



1. Servers
2. HP P4000
3. Storage cluster
4. GigE trunk
5. Active path
6. Passive path

Figure 26 (page 60) illustrates the Active-Passive configuration in a four-switch topology.

How link aggregation dynamic mode bonding works

Link Aggregation Dynamic Mode allows the storage system to use both interfaces simultaneously for data transfer. Both interfaces have an active status. If the interface link to one NIC goes offline, the other interface continues operating. Using both NICs also increases network bandwidth.

Requirements for link aggregation dynamic mode

To configure Link Aggregation Dynamic Mode:

- Both NICs should be enabled.
- NICs must be configured to the same subnet.
- NICs must be connected to a single switch that is LACP-capable and supports 802.3ad link aggregation. If the storage system is directly connected to a server, then the server must support 802.3ad link aggregation.

Which physical interface is preferred

Because the logical interface uses both NICs simultaneously for data transfer, neither of the NICs in an aggregation bond is designated as preferred.

Which physical interface is active

When the Link Aggregation Dynamic Mode bond is created, if both NICs are plugged in, both interfaces are active. If one interface fails, the other interface continues operating. For example, suppose Eth0 and Eth1 are bonded in a Link Aggregation Dynamic Mode bond. If Eth0 fails, then Eth1 remains active.

Once the link is fixed and Eth0 is operational, it becomes active again. Eth1 remains active.

Table 22 Link aggregation dynamic mode failover scenario and corresponding NIC status

Example failover scenario	NIC status
1. Link Aggregation Dynamic Mode bond0 is created. Eth0 and Eth1 are both active.	<ul style="list-style-type: none"> Bond0 is the master logical interface. Eth0 is Active. Eth1 is Active.
2. Eth0 interface fails. Because Link Aggregation Dynamic Mode is configured, Eth1 continues operating.	<ul style="list-style-type: none"> Eth0 status becomes Passive (Failed). Eth1 status remains Active.
3. Eth0 link failure is repaired.	<ul style="list-style-type: none"> Eth0 resumes Active status. Eth1 remains Active.

Summary of NIC states during failover

Table 23 (page 61) shows the states of Eth0 and Eth1 when configured for Link Aggregation Dynamic Mode.


Table 23 NIC status during failover with link aggregation dynamic mode

Failover status	Status of Eth0	Status of Eth1
Normal Operation	Preferred: No Status: Active Data Transfer: Yes	Preferred: No Status: Active Data Transfer: Yes
Eth0 Fails, Data Transfer Fails Over to Eth1	Preferred: No Status: Passive (Failed) Data Transfer: No	Preferred: No Status: Active Data Transfer: Yes
Eth0 Restored	Preferred: No Status: Active Data Transfer: Yes	Preferred: No Status: Active Data Transfer: Yes

Example network cabling topologies with link aggregation dynamic mode

A simple network configuration using Link Aggregation Dynamic Mode in a high-availability environment is illustrated in Figure 27 (page 62).

Figure 27 Link aggregation dynamic mode in a single-switch topology



1. Servers
2. HP P4000
3. Storage cluster

How Adaptive Load Balancing works

Adaptive Load Balancing allows the storage system to use both interfaces simultaneously for data transfer. Both interfaces have an active status. If the interface link to one NIC goes offline, the other interface continues operating. Using both NICs also increases network bandwidth.

Requirements for Adaptive Load Balancing

To configure Adaptive Load Balancing:

- Both NICs must be enabled.
- NICs must be configured to the same subnet.
- NICs can be connected to separate switches.

Which physical interface is preferred

Because the logical interface uses both NICs for data transfer, neither of the NICs in an Adaptive Load Balancing bond is designated as preferred.

Which physical interface is active

When the Adaptive Load Balancing bond is created, if both NICs are plugged in, both interfaces are active. If one interface fails, the other interface continues operating. For example, suppose Motherboard:Port1 and Motherboard:Port2 are bonded in an Adaptive Load Balancing bond. If Motherboard:Port1 fails, then Motherboard:Port2 remains active.

Once the link is fixed and Motherboard:Port1 is operational, it becomes active again. Motherboard:Port2 remains active.

Table 24 Example Adaptive Load Balancing failover scenario and corresponding NIC status

Example failover scenario	NIC status
1. Adaptive Load Balancing bond0 is created. Motherboard:Port1 and Motherboard:Port2 are both active.	<ul style="list-style-type: none"> • Bond0 is the master logical interface. • Motherboard:Port1 is Active. • Motherboard:Port2 is Active.
2. Motherboard:Port1 interface fails. Because Adaptive Load Balancing is configured, Motherboard:Port2 continues operating.	<ul style="list-style-type: none"> • Motherboard:Port1 status becomes Passive (Failed). • Motherboard:Port2 status remains Active.
3. Motherboard:Port1 link failure is repaired.	<ul style="list-style-type: none"> • Motherboard:Port1 resumes Active status. • Motherboard:Port2 remains Active.

Summary of NIC states during failover

Table 25 (page 63) shows the states of Motherboard:Port1 and Motherboard:Port2 when configured for Adaptive Load Balancing.


Table 25 NIC status during failover with Adaptive Load Balancing

Failover status	Status of Motherboard:Port1	Status of Motherboard:Port2
Normal Operation	Preferred: No Status: Active Data Transfer: Yes	Preferred: No Status: Active Data Transfer: Yes
Motherboard:Port1 Fails, Data Transfer Fails Over to Motherboard:Port2	Preferred: No Status: Passive (Failed) Data Transfer: No	Preferred: No Status: Active Data Transfer: Yes
Motherboard:Port1 Restored	Preferred: No Status: Active Data Transfer: Yes	Preferred: No Status: Active Data Transfer: Yes

Example network cabling topologies with Adaptive Load Balancing

A simple network configuration using Adaptive Load Balancing in a high availability environment is illustrated in Figure 28 (page 64).

Figure 28 Adaptive Load Balancing in a two-switch topology



1. Servers
2. HP P4000
3. Storage cluster
4. GigE trunk

Creating a NIC bond

Follow these guidelines when creating NIC bonds:

Prerequisites

Verify that the speed, duplex, flow control, and frame size are all set properly on both interfaces that are being bonded. These settings cannot be changed on a bonded interface or on either of the supporting interfaces.

For detailed instructions about properly configuring these settings, see [“Managing settings on network interfaces”](#) (page 50).

Bond guidelines

- Create a bond on a storage system before you add the storage system to a management group.
- Create bonds of two interfaces.
- An interface can only be in one bond.
- Record the configuration information of each interface before you create the bond. Then, if you delete the bond, you can return to the original configuration if desired.
 - When you delete an Active-Passive bond, the preferred interface assumes the IP address and configuration of the deleted logical interface.
 - When you delete a Link Aggregation Dynamic Mode or an Adaptive Load Balancing bond, one of the interfaces retains the IP address of the deleted logical interface. The IP address of the other interface is set to 0.0.0.0.

- Ensure that the bond has a static IP address for the logical bond interface. The default values for the IP address, subnet mask and default gateway are those of one of the physical interfaces.
- Verify on the Communication tab that the SAN/iQ interface is communicating with the bonded interface.

⚠ CAUTION: To ensure that the bond works correctly, you should configure it as follows:

- Create the bond on the storage system before you add it to a management group.
- Verify that the bond is created.

If you create the bond on the storage system after it is in a management group, and if it does not work correctly, you might

- Lose the storage system from the network
- Lose quorum in the management group for a while.

See “Deleting a NIC bond” (page 237) for information about deleting NIC bonds using the Configuration Interface.

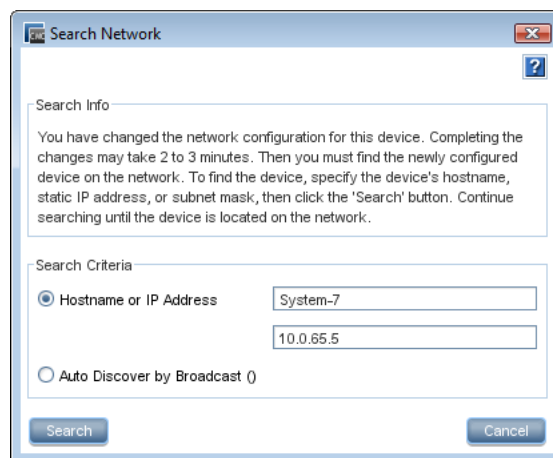
Creating the bond

1. Log in to the storage system.
2. Select the **TCP/IP** category from the tree.
3. On the TCP/IP tab, select both NICs to bond.
4. Click **TCP/IP Tasks**, and select **New Bond**.
5. Select a bond type from the drop-down list.
6. Enter an IP address for the bond or accept the default.
7. Enter the Subnet mask.
8. (Optional) Enter the default gateway.
9. Click **OK**.

NOTE: The storage system drops off the network while the bonding takes place. The changes may take 2 to 3 minutes, during which time you cannot find or access the storage system.

10. Click **OK** to confirm the TCP/IP changes.
A message opens, prompting you to search for the bonded storage system on the network.

Figure 29 Searching for the bonded storage system on the network



11. Search for the storage system by Host Name or IP address, or by Subnet/mask.

NOTE: Because it can take a few minutes for the storage system to re-initialize, the search may fail the first time. If the search fails, wait a minute or two and choose Try Again on the Network Search Failed message.

12. Verify the new bond interface.

Figure 30 Viewing a new Active-Passive bond

TCP/IP TCP Status DNS Routing Communication							
Name	Descript...	MAC	Mode	IP Addre...	Subnet ...	Gateway	Type
bond0	Logical Fa...	00:25:B3:...	Static IP	10.0.65.5	255.255.0.0	10.0.255.2...	Active - P...
+ Mother...	Intel Corp...	00:25:B3:...	Slave	10.0.65.5	255.255.0.0	10.0.255.2...	NIC
+ Mother...	Intel Corp...	00:25:B3:...	Slave	10.0.65.5	255.255.0.0	10.0.255.2...	NIC

1. Bonded logical network interface
2. Physical interfaces shown as slaves

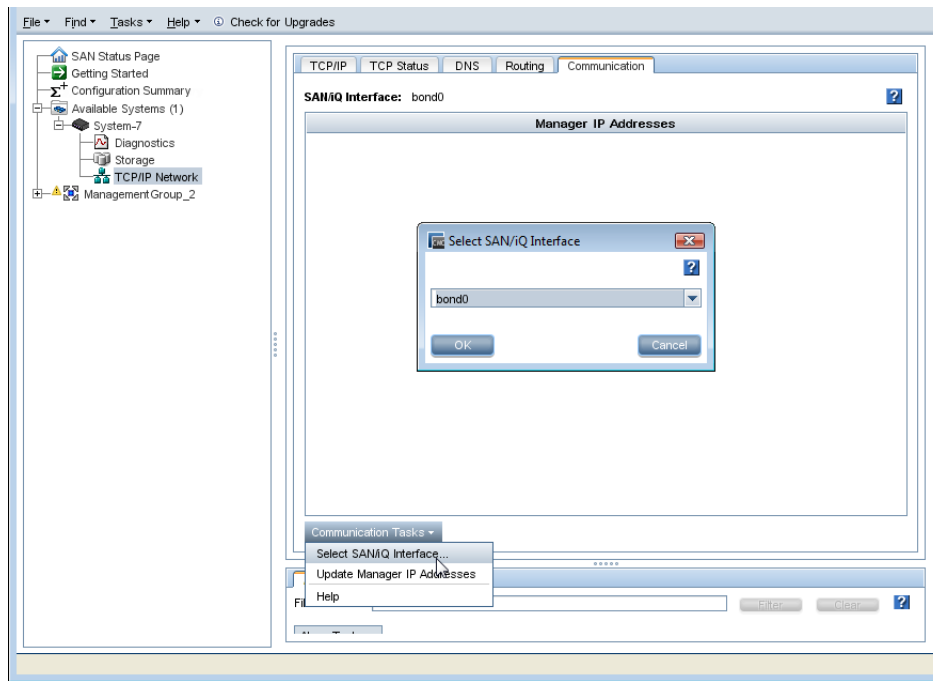
The bond interface shows as “bond0” and has a static IP address. The two physical NICs now show as slaves in the Mode column.

13. (Optional, for Active-Passive bonds only) To change which interface is the preferred interface in an Active-Passive bond, on the TCP Status tab select one of the NICs in the bond, and click **Set Preferred**.

Verify communication setting for new bond

1. Select a storage system, and open the tree below it.
2. Select the **TCP/IP Network** category, and click the **Communication** tab.

Figure 31 Verifying interface used for SAN/iQ communication



3. Verify that the SAN/iQ communication port is correct.

Viewing the status of a NIC bond

You can view the status of the interfaces on the TCP Status tab. Notice that in the Active-Passive bond, one of the NICs is the preferred NIC. In both the Link Aggregation Dynamic Mode bond and the Adaptive Load Balancing bond, neither physical interface is preferred.

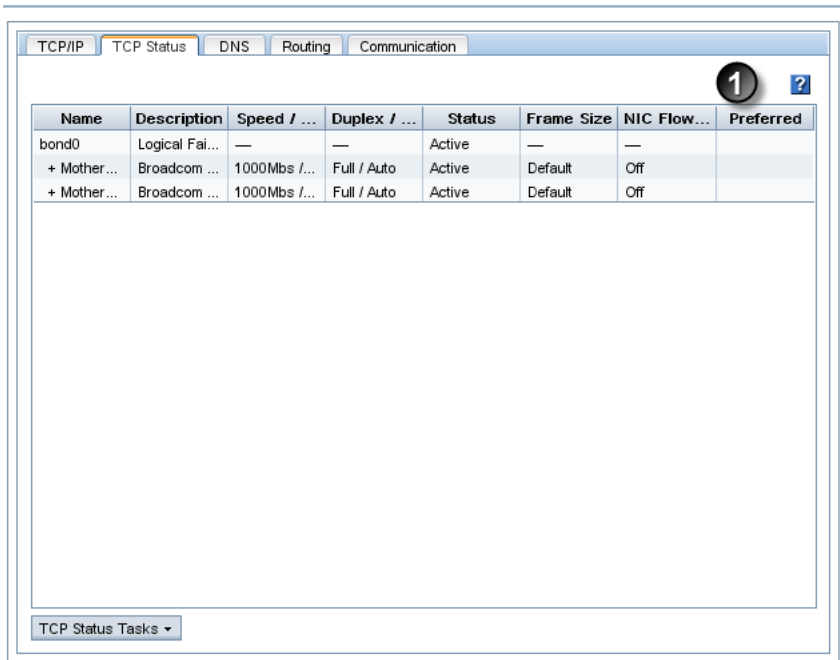
Figure 32 (page 67) shows the status of interfaces in an Active-Passive bond. Figure 33 (page 68) shows the status of interfaces in a Link Aggregation Dynamic Mode bond.

Figure 32 Viewing the status of an Active-Passive bond

TCP/IP TCP Status Routing Communication							
SAN/iQ Interface: bond0							
Manager IP Addresses							
Select SAN/iQ Interface							
bond0							
OK Cancel							
Communication Tasks							
Select SAN/iQ Interface... Update Manager IP Addresses Help							
TCP Status Tasks							
Name	Description	Speed / ...	Duplex / ...	Status	Frame Size	NIC Flow...	Preferred
bond0	Logical Fai...	—	—	Active	—	—	
+ Mother...	Intel Corpo...	1000Mbps /...	Full / Auto	Passive (...)	Default	Off	
+ Mother...	Intel Corpo...	1000Mbps /...	Full / Auto	Active	Default	Off	Yes

1. Preferred interface

Figure 33 Viewing the status of a link aggregation dynamic mode bond



1. Neither interface is preferred

NOTE: If the bonded NIC experiences rapid, sequential Ethernet failures, the CMC may display the storage system as failed (flashing red) and access to data on that storage system fails. However, as soon as the Ethernet connection is reestablished, the storage system and the CMC display the correct information.

Deleting a NIC bond

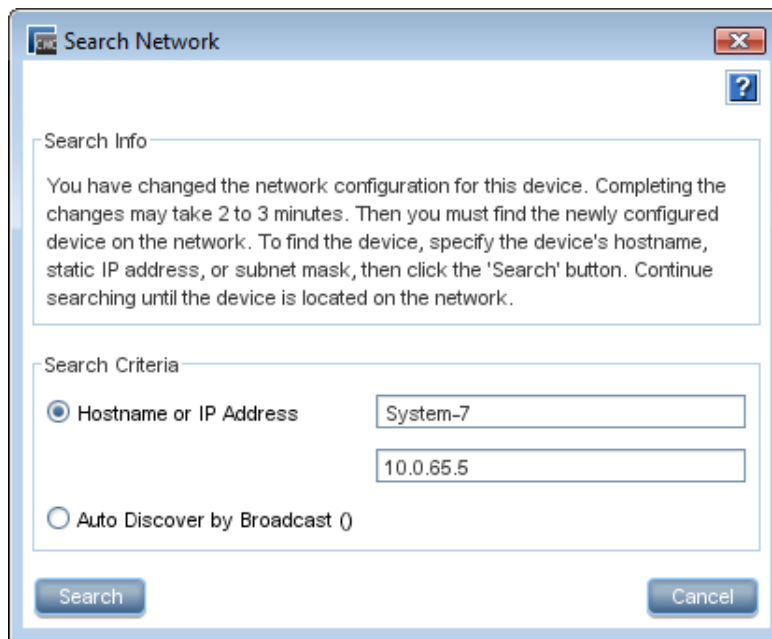
When you delete an Active-Passive bond, the preferred interface assumes the IP address and configuration of the deleted logical interface. The other NIC is disabled, and its IP address is set to 0.0.0.0.

When you delete either a Link Aggregation Dynamic Mode or an Adaptive Load Balancing bond, one of the active interfaces in the bond retains the IP address of the deleted logical interface. The other NIC is disabled, and its IP address is set to 0.0.0.0.

1. Log in to the storage system, and expand the tree.
2. Select the **TCP/IP** category from the tree.
3. On the TCP/IP tab, select the bond interface or physical bond to delete.
4. Click **TCP/IP Tasks**, and select **Delete Bond**.

Because the IP addresses changes, the Search Network window opens.

Figure 34 Searching for the unbonded storage system on the network



5. Search for the storage system by Host Name or IP Address or Subnet/Mask.

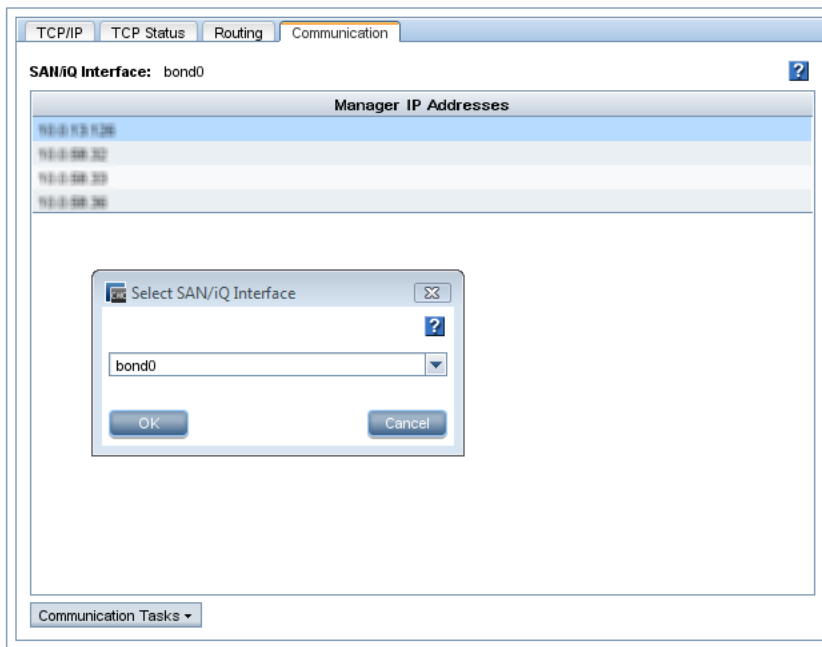
NOTE: Because it can take a few minutes for the storage system to re-initialize, the search may fail the first time. If the search fails, wait a minute or two and choose Try Again on the Network Search Failed message.

You can also use the Configuration Interface to delete a NIC bond. See [“Deleting a NIC bond”](#) (page 237).

Verify NIC settings and communication settings after deleting a bond

1. Select a storage system, and open the tree below it.
2. Select the **TCP/IP Network** category.
3. Check the interfaces on the TCP/IP tab, and reconfigure them if necessary.
After deleting a bond, the interfaces may not have the correct IP addresses, or one interface may be disabled.
4. Click the **Communication** tab.

Figure 35 Verifying interface used for SAN/iQ communication



5. Verify that the SAN/iQ communication port is correct.

Disabling a network interface

You can disable the network interfaces on the storage system.

- You can only disable top-level interfaces. This includes bonded interfaces and NICs that are not part of bonded interfaces.
- To ensure that you always have access to the storage system, do not disable the last interface. If you want to disable the last interface, first enable another interface.

CAUTION: If you disable an interface, be sure you enable another interface first. That way you always have access to the storage system. If you disable all the interfaces, you must reconfigure at least one interface using the Configuration Interface to access the storage system. See [“Configuring a network connection” \(page 236\)](#).

To disable a network interface

1. Log in to the storage system, and open the tree.
2. Select the **TCP/IP Network** category.
3. Select from the list on the TCP/IP tab window the interface to disable.
4. Click **TCP/IP Tasks**, and select **Edit**.
5. Click **Disable Interface**.
6. Click **OK**.

A confirmation message opens. If you are disabling the only interface, the message warns that the storage system may be inaccessible if you continue.

7. Click **OK**.

If the storage system for which you are disabling the interface is a manager in a management group, a window opens which displays all the IP addresses of the managers in the management group and a reminder to reconfigure the application servers that are affected by the update.

Configuring a disabled interface

If one interface is still connected to the storage system but another interface is disconnected, you can reconnect to the second interface using the CMC. See [“Configuring the IP address manually” \(page 54\)](#).

If both interfaces to the storage system are disconnected, you must attach a terminal, or PC or laptop to the storage system with a null modem cable and configure at least one interface using the Configuration Interface. See [“Configuring a network connection” \(page 236\)](#).

Using a DNS server

You configure DNS at the management group level for all storage systems in the management group. Use the DNS tab for the management group to change the settings.

The storage system can use a DNS server to resolve host names. For example, if you enter a host name to specify an NTP time server, the storage system will use DNS to resolve the host name to its IP address. For example, the time server in Boulder, Colorado has a host name of `time.nist.gov`. DNS resolves this host name to its IP address of 192.43.244.18.

DNS and DHCP

If you configure the storage system to use DHCP to obtain an IP address, and if the DHCP server is configured to provide the IP addresses of the DNS servers, then a maximum of three DNS servers will automatically be added to the management group. These DNS servers are listed as IP addresses in the management group on the DNS tab. You can remove these DNS servers, but storage systems will not be able to resolve host names until you enter a new DNS server.

DNS and static IP addresses

If you assigned a static IP address to the storage system and you want the storage system to recognize host names, you must manually add a DNS server to the management group DNS tab.

NOTE: If you initially set up the storage system to use DHCP, and then change the configuration to use a static IP address, the DNS server provided by DHCP will remain on the DNS tab. You can remove or change this DNS server.

Getting there

1. In the navigation window, select a management group and log in.
2. Select the **DNS** tab.

Adding the DNS domain name

Add the name of the DNS domain in which the management group resides.

1. Click **DNS Tasks**, and select **Edit DNS Domain Name**.
2. Enter the DNS domain name.
3. Click **OK** when you are finished.

Adding the DNS server

Add up to three DNS servers for use with the management group.

1. Click **DNS Tasks**, and select **Edit DNS Servers**.
2. Click **Add**, and enter the IP address for the DNS server.
3. Click **OK**.
4. Repeat steps 1 through 3 to add up to three servers.

5. Use the arrows on the Edit DNS Servers window to order the servers.
The servers will be accessed in the order they appear in the list.
6. Click **OK** when you are finished.

Adding domain names to the DNS suffixes

Add up to six domain names to the DNS suffix list (also known as the look-up zone). The storage system searches the suffixes first and then uses the DNS server to resolve host names.

1. On the DNS tab, click **DNS Tasks**, and select **Edit DNS Suffixes**.
2. Click **Add** to display the Add DNS Suffixes window.
3. Enter the DNS suffix name. Use the domain name format.
4. Click **OK**.
5. Repeat steps 1 through 4 to add up to six domain names.
6. Click **OK** when you are finished.

Editing a DNS server

Change the IP address for a DNS Server in the list.

1. In the navigation window, select a management group and log in.
2. Select the **DNS** tab.
3. Select the server to edit.
4. Click **DNS Tasks**, and select **Edit DNS Servers**.
5. Select the server again, and click **Edit**.
6. Enter the new IP address for the DNS server, and click **OK**.

Editing a domain name in the DNS suffixes list

Change a domain name of a management group.

1. In the navigation window, select a management group and log in.
2. Select the **DNS** tab.
3. Click **DNS Tasks**, and select **Edit DNS Domain Name**.
4. Enter the change to the domain name.
5. Click **OK**.

Removing a DNS server

Remove a DNS server from the list.

1. In the navigation window, select a management group and log in.
2. Select the **DNS** tab.
3. Select the server to remove from the DNS Servers list.
4. Click **DNS Tasks**, and select **Edit DNS Servers**.
5. Select the name again in the Edit DNS Servers window.
6. Click **Remove**.
7. Click **OK** to remove the DNS server from the list.

Removing a domain suffix from the DNS suffixes list

1. In the navigation window, select a management group and log in.
2. Select the **DNS** tab.
3. Select the suffix to remove.
4. Click **DNS Tasks**, and select **Edit DNS Suffixes**.
5. Select the name again in the Edit DNS Suffixes window.

6. Click **Remove**.
7. Click **OK** to remove the DNS suffix from the list.

Setting up routing

The Routing tab displays the routing table. You can specify static routes and/or a default route.

NOTE: If you specify a default route here, it will not survive a reboot or shutdown of the storage system. To create a route that will survive a storage system reboot or shut down, you must enter a default gateway on the TCP/IP tab. See [“Configuring the IP address manually” \(page 54\)](#).

Information for each route listed includes the device, the network, gateway, mask, and flags.

Adding routing information

1. In the navigation window, select a storage system and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **Routing** tab.
4. Click **Routing Tasks**, and select **Edit Routing Information**.
5. Click **Add**.
6. Select the port to use for routing in the Device list.
7. Enter the IP address portion of the network address in the Net field.
8. Enter the IP address of the router in the Gateway field.
9. Select the netmask.
10. Click **OK**.
11. Use the arrows on the routing table panel to order devices according to the configuration of your network.

The storage system attempts to use the routes in the order in which they are listed.

Editing routing information

You can only edit optional routes you have added.

1. In the navigation window, select a storage system, and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **Routing** tab.
4. On the Routing tab, select the optional route to change.
5. Click **Routing Tasks**, and select **Edit Routing Information**.
6. Select a Route, and click **Edit**.
7. Change the relevant information.
8. Click **OK**.

Deleting routing information

You can only delete optional routes you have added.

1. In the navigation window, select a storage system, and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **Routing** tab.
4. On the Routing tab, select the optional route to delete.
5. Click **Routing Tasks**, and select **Edit Routing Information**.
6. Select the routing information row to delete.
7. Click **Delete**.
8. Click **OK** on the confirmation message.

Configuring storage system communication

Use the Communication tab to configure the network interface used by the storage system to communicate with other storage systems on the network and to update the list of managers that the storage system can communicate with.

Selecting the interface used by the SAN/iQ software

The SAN/iQ software uses one network interface for communication with other storage systems on the network. In order for clustering to work correctly, the SAN/iQ software communication interface must be designated on each storage system. The interface can be

- A single NIC that is not part of a bond
- A bonded interface consisting of two bonded NICs

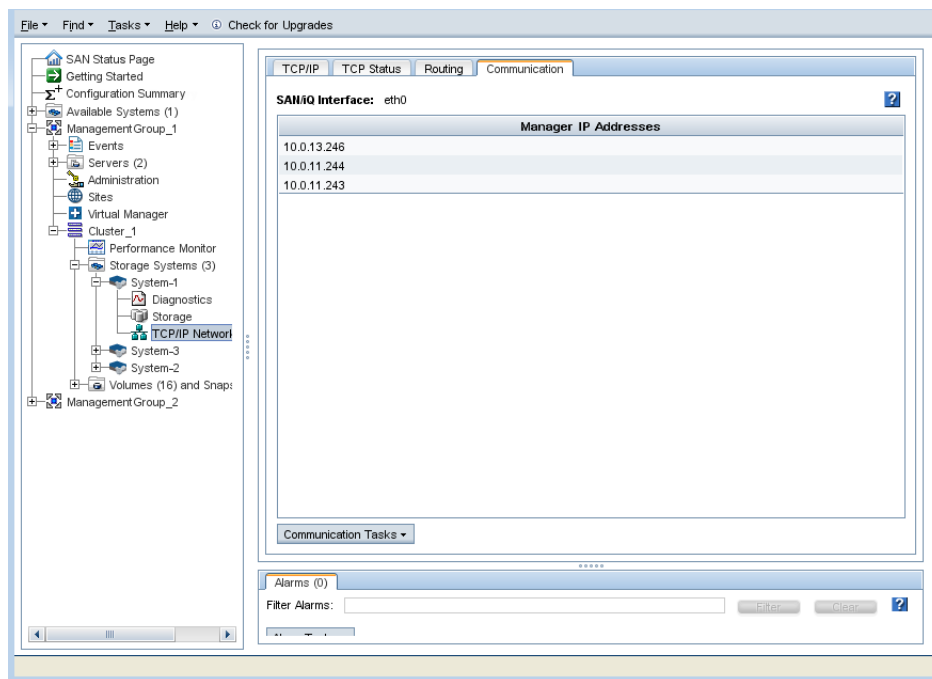
NOTE: Only NICs that are in the Active or Passive (Ready) state can be designated as the communication interface. You cannot make a disabled NIC the communication interface.

When you initially set up a storage system using the Configuration Interface, the first interface that you configure becomes the interface used for the SAN/iQ software communication.

To select a different communication interface:

1. In the navigation window, select the storage system, and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **Communication** tab to bring that window to the front.

Figure 36 Selecting the SAN/iQ software network interface and updating the list of managers



4. Select an IP address from the list of Manager IP Addresses.
5. Click **Communication Tasks**, and select **Select SAN/iQ Address**.
6. Select an Ethernet port for this address.
7. Click **OK**.

Now, this storage system connects to the IP address through the Ethernet port you selected.

Updating the list of manager IP addresses

Update the list of manager IP addresses to ensure that a manager running on this storage system is communicating correctly with all managers in the management group.

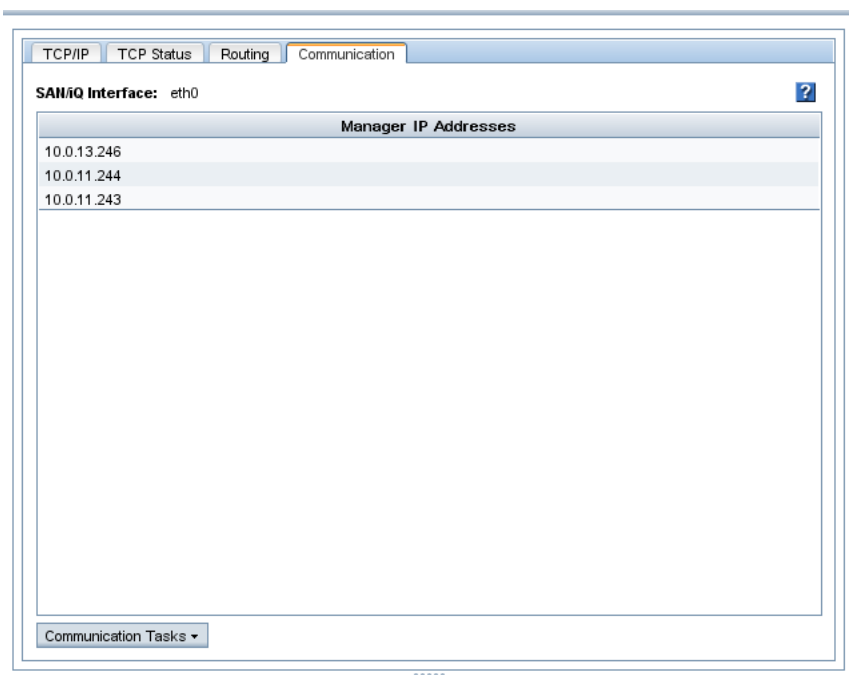
Requirements

Each time you update the list of managers, you must reconfigure application servers that use the management group to which this storage system belongs. Only update the list mode if you have reason to believe that there is a problem with the communication between the other managers in the group and the manager on this storage system.

To update the list of IP addresses

1. In the navigation window, select a storage system, and log in.
2. Open the tree, and select the **TCP/IP Network** category.
3. Select the **Communication** tab.

Figure 37 Viewing the list of manager IP addresses



4. Click **Communication Tasks**, and select **Update Communications List**.

The list is updated with the current storage system in the management group and a list of IPs with every manager's enabled network interfaces.

A window opens which displays the manager IP addresses in the management group and a reminder to reconfigure the application servers that are affected by the update.

5 Setting the date and time

The storage systems within management groups use the date and time settings to create a time stamp when data is stored. You set the time zone and the date and time in the management group, and the storage systems inherit those management group settings.

- **Using network time protocol**

Configure the storage system to use a time service, either external or internal to your network.

- **Setting the time zone**

Set the time zone for the storage system. The time zone controls the time stamp on volumes and snapshots.

If you use NTP, decide what time zone you will use. You can use either GMT on all of your management groups, or you can set each management group to its local time.

If you do not set the time zone for each management group, the management group uses the GMT time zone, whether or not you use NTP.

- **Setting date and time**

Set the date and time on the management group(s) if not using an NTP time service.

Management group time

When you create a management group, you set the time zone and the date and time while going through the Management Groups, Clusters and Volumes wizard. This ensures that all the storage systems in the management group have the same time setting.

Getting there

1. In the network window, select a management group and log in.
2. Click the **Time** tab.

Refreshing the management group time

Use **Refresh All** to update the view of the time on all storage systems in the management group. This view is not updated automatically, so you must refresh the view to verify that the time settings on the storage systems are what you expect.

1. Select a management group.
2. Click the **Time** tab.
3. Select **Time Tasks**, and select **Refresh All**.

After processing, all storage systems display the current time.

Using NTP

Network time protocol servers (NTP) can manage the time for the management group instead of using the local system time. NTP updates occur at five-minute intervals. If you do not set the time zone for the management group, it uses GMT.

NOTE: When using a Windows server as an external time source for an storage system, you must configure W32Time (the Windows Time service) to also use an external time source. The storage system does not recognize the Windows server as an NTP server if W32Time is configured to use an internal hardware clock.

1. Click **Time Tasks**, and select **Add NTP Server**.
2. Enter the IP address of the NTP server you want to use.
3. Decide whether you want this NTP server to be designated preferred or not preferred.

NOTE: A **preferred** NTP server is one that is more reliable, such as a server that is on a local network. An NTP server on a local network would have a reliable and fast connection to the storage system. **Not preferred** designates an NTP server to be used as a backup if a preferred NTP server is not available. An NTP server that is *not* preferred might be one that is located elsewhere or has a less reliable connection.

4. Click **OK**.

The NTP server is added to the list on the NTP tab.

The NTP servers are accessed in the order you add them, and preferred servers are accessed before non-preferred servers. The first server you add, if it is marked preferred, has the highest order of precedence. The second server you add takes over as a time server if the preferred server fails.

Editing NTP servers

Change whether an NTP server is preferred or not.

1. Select an NTP server in the list.
2. Click **Time Tasks**, and select **Edit NTP Server**.
3. Change the preference of the NTP server.
4. Click **OK**.

The list of NTP servers displays the changed NTP server in the list.

NOTE: To change the IP address of an NTP server, you must remove the server no longer in use and add a new NTP server.

Deleting an NTP server

You may need to delete an NTP server:

- If the IP address of that server becomes invalid
- If you no longer want to use that server
- If you want to change the order of servers in the list

Delete an NTP server

1. Select an NTP server in the list on the Time tab window.
2. Click **Time Tasks**, and select **Delete NTP Server**.
3. Click **OK** on the confirmation window.

The list of NTP servers refreshes the list of available servers.

Changing the order of NTP servers

The window displays the NTP servers in the order you added them.

The server you added first is the one accessed first when time needs to be established. If this NTP server is not available for some reason, the next NTP server that was added, and is preferred, is used for time serving.

To change the order of access for time servers

1. Delete the server whose place in the list you want to change.
2. Add that same server back into the list.

It is placed at the bottom of the list, and is the last to be accessed.

Editing the date and time

You initially set the date and time when you create the management group using the Management Groups, Clusters and Volumes wizard. If necessary, you can edit these settings later.

1. Select the management group.
2. Select the **Time** tab to bring it to the front.
3. Click **Time Tasks**, and select **Edit Date, Time, Time Zone**.
4. Change the date and time to the correct date and time for that time zone.
 - In the Date group box, set the year, month, and day.
 - In the Time group box, highlight a portion of the time and increase or decrease it with the arrows. You may also enter in the time directly.
 - Select a time zone for the Time Zone drop-down list.

NOTE: If you use an NTP server, you have the option of setting the time zone only.

5. Click **OK**.

A warning message informs you that there may be a slight time lag for a reset to take effect.

6. Click **OK**.

Editing the time zone only

You initially set the time zone when you create the management group. You can change the time zone later, if necessary.

If you do not set the time zone for each management group, the management group uses GMT, whether or not you use NTP. Files display the time stamp according to this local time zone.

1. Click **Time Tasks**, and select **Edit Time Zone**.
2. From the drop-down list, select the time zone in which this management group resides.
3. Click **OK**.

Note the change in the Time column of the **Time** tab window.

6 Administrative users and groups

When you create a management group, the SAN/iQ software configures two default administrative groups and one default administrative user. You can add, edit, and delete additional administrative users and groups. All administrative users and groups are managed at the management group level.

Getting there

In the navigation window, log in to the management group, and select the **Administration** node.

Managing administrative users

When you create a management group, one default administrative user is created. Use the default user and/or create new ones.

Default administrative user

The user who is created when you create a management group becomes a member of the Full Administrator group by default.

Adding a new administrative user

Add administrative users as necessary to provide access to the management functions of the SAN/iQ software.

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **New User**.
3. Enter a **User Name** and **Description**.
4. Enter a password, and confirm that password.
5. Click **Add** in the Member Groups section.
6. Select one or more groups to which you want the new user to belong.
7. Click **OK**.
8. Click **OK** to finish adding the administrative user.

Editing administrative users

Each management group has an administration node in the tree below it. You can add, edit, and remove administrative users here. Editing administrative users includes changing passwords and group memberships of administrative users.

Changing a user's description

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit User**.
3. Change the User Description as necessary.
4. Click **OK** to finish.

Changing a user's password

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit User**.
3. Enter a new password, and confirm that password.
4. Click **OK** to finish.

Adding group membership to a user

1. Log in to the management group, and select the **Administration** node.
2. Select a user in the Users table.
3. Click **Administration Tasks** in the tab window, and select **Edit User**.
4. Click **Add** in the Member Groups section.
5. Select the groups to which to add the new user.
6. Click **OK**.
7. Click **OK** to finish editing the administrative user.

Removing group membership from a user

1. Log in to the management group, and select the **Administration** node.
2. Select a user in the Users table.
3. Click **Administration Tasks** in the tab window, and select **Edit User**.
4. In the Member Groups section, select the group from which to remove the user.
5. Click **Remove**.
6. Click **OK** to finish editing the administrative user.

Deleting an administrative user

1. Log in to the management group, and select the **Administration** node.
2. Select a user in the **Users** table.
3. Click **Administration Tasks** in the tab window, and select **Delete User**.
4. Click **OK**.

NOTE: If you delete an administrative user, that user is automatically removed from any administrative groups.

Managing administrative groups

When you create a management group, two default administrative groups are created. Use these groups and/or create new ones.

Default administrative groups

The two default administrative groups and the permissions granted to those groups are listed in [Table 26 \(page 80\)](#). Users assigned to either of these groups assume the privileges associated with that group.

Table 26 Using default administrative groups

Name of group	Management capabilities assigned to group
Full_Administrator	Manage all functions (read/write access to all functions)
View_Only_Administrator	View-only capability to all functions (read only)

Administrative groups can have:

- Different levels of access to the storage system, such as read/write
- Access to different management capabilities for the SAN, such as configuring network capabilities

Adding administrative groups

When you create a group, you also set the management permissions for the users assigned to that group. The default setting for a new group is Read Only for each category.

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **New Group**.
3. Enter a Group Name and optional Description.
4. Select the permission level for each function for the group you are creating. See [Table 27](#) (page 81).

Table 27 Descriptions of group permissions

Management area	Activities controlled by this area
Change Password	User can change other administrative users' passwords.
Management Groups, RAID, Drive Hot Swap	User can set the RAID configuration for the storage system. Shut down disks, restart RAID, and hot-swap disks. Create management groups.
Network	User can choose type of network connection, set the time and time zone for the management group, identify the Domain Name Server, and use SNMP.
Storage System Administration, Boot, Upgrade	User can add administrators and upgrade the SAN/iQ software.
System and Disk Report	User can view reports about the status of the storage system.

What the permission levels mean

- **Read Only**—User can only view the information about these functions.
 - **Read-Modify**—User can view and modify existing settings for these functions.
 - **Full**—User can perform all actions (view, modify, add new, delete) in all functions.
1. Add a user to the group.
 - Click **Add** in the Users section.
 - Select one or more users to add to the group.
 - Click **Add**.
 2. Click **OK** to finish creating a new group.

Editing administrative groups

Each management group has an administration node in the tree below it. You can add, edit, and remove administrative groups here. Editing an administrative group includes changing the description, permissions, and users for the group.

Change the description of a group

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit Group**.
3. Change the Description as necessary.
4. Click **OK** to finish.

Changing administrative group permissions

Change the management capabilities available to members of a group.

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit Group**.
Administrative groups can have:

- Different levels of access to the storage system, such as read/write
- Access to different management capabilities for the storage system, such as creating volumes

When you create a group, you also set the management capabilities available to members of a group. The default setting for a new group is Read Only for each category.

3. Click the permission level for each function for the group you are creating.
See [“Descriptions of group permissions” \(page 81\)](#) for a description of the permission levels.
4. Click **OK** to finish.

Adding users to an existing group

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit Group**.
3. Click **Add** in the Users section.
4. Select one or more users from list of administrative users to add to the group.
5. Click **Add**.
6. Click **OK** to finish creating a new group.

Removing users from a group

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Edit Group**.
3. Select one or more users to remove from the group.
4. Click **Remove**.
5. Click **OK** to finish.

Deleting administrative groups

Delete all users from a group before you delete the group.

1. Log in to the management group, and select the **Administration** node.
2. Click **Administration Tasks** in the tab window, and select **Delete Group**.
3. Click **OK** on the confirmation window.
4. Click **OK** to finish.

7 Monitoring the SAN

Monitor the SAN for usage; to ensure that best practices are followed when changes are made, such as adding additional storage systems to clusters; and maintain the overall health of the SAN. Tools for monitoring the SAN include the SAN Status Page, the Configuration Summary and the Best Practice table, the Alarms and Events features, including customized notification methods, and diagnostic tests and log files available for the storage systems.

Monitoring SAN status

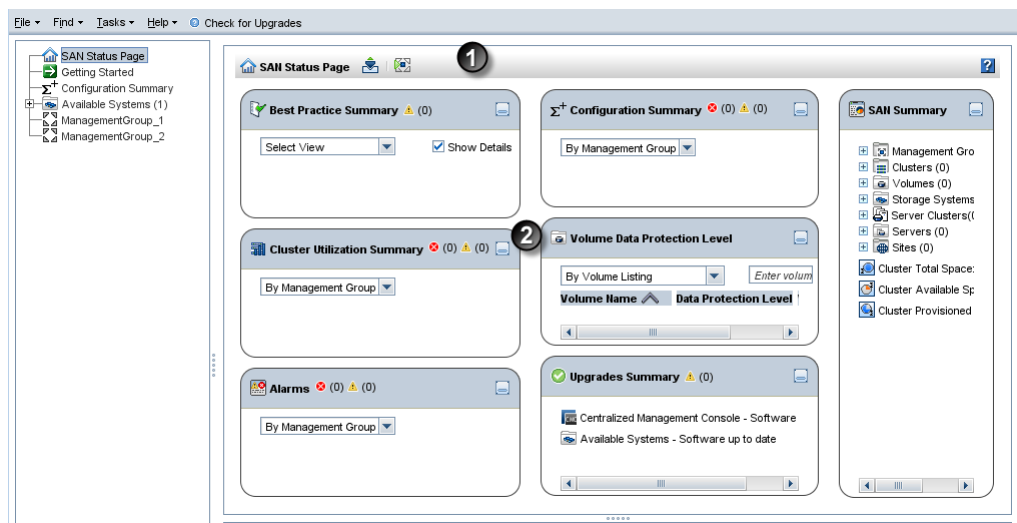
After you have created at least one management group in the CMC, the CMC searches for storage systems, and then opens to the SAN Status Page, shown in [Figure 38 \(page 83\)](#). The SAN Status Page brings together multiple categories of information to provide one location from which to monitor multiple features of the SAN at one time. Use the SAN Status Page to monitor the following items:

- **Best Practices**—Ensure that the SAN configurations meet best practice recommendations for data protection, SAN and data availability, and network connectivity, failover, and performance.
- **Configurations**—Monitor SAN configurations to ensure optimum capacity management, performance, availability, and ease of management.
- **Utilization**—Track cluster use, space on the SAN, and summary information about configurations that affect performance and scalability.
- **Data protection**—Ensure that the SAN is configured for ongoing maximum data protection while scaling capacity and performing system maintenance.
- **Alarms**—Manage the SAN health by easily tracking issues that may arise.

Configuring the SAN Status Page

Customize the information displayed on the Status Page.

Figure 38 SAN Status Page



1. Status Page toolbar
2. Content panes

Customizing the Status Page content

The SAN Status Page includes seven categories of information, displayed in content panes. You can customize the layout of the content panes, as well customize the information displayed in each content pane.

- **Best Practice Summary**—includes all the best practice items that appear on the Configuration Summary node. View the information by management group. Select **Show Details** to turn on and off the pop-up descriptions of the best practices.
- **Cluster Utilization Summary**—displays cluster utilization information by management group or by utilization percentage by cluster.
- **Alarms**—summarizes alarms by management group
- **Configuration Summary**—displays configuration summary information about the SAN by management group. Additional views on the Status Page content pane include summary by configuration type and summary by optimum number. Optimum number refers to number of items by category, displayed in descending order, according to the relative number for a category. For example, six storage systems in a cluster are closer to the optimum number for that cluster (10) than they are for the number of storage systems in a management group (20). Therefore, as displayed in the Optimum Number view, the six storage systems in the cluster will be above those same six storage systems in the management group.
- **Volume Data Protection Level**—lists volumes by their Network RAID level. Network RAID 0 provides no data protection, while the other Network RAID levels offer varying levels of data protection. Filter volumes by name to control the view by grouping volumes that share naming conventions.
- **Upgrades Summary**—identifies which components of the SAN are up-to-date and which need upgrading.
- **SAN Summary**—provides a comprehensive overview of the entire SAN, including aggregated cluster space.

Customizing the Status Page layout

Customize the layout of the SAN Status Page to highlight the information most important to you. All customizations are retained when the CMC is closed and restarted. If the Status Page is undocked when the CMC is closed, it will be undocked when the CMC is started again.

Drag-and-drop content panes to change their position on the page. The layout is three columns by default. Drag a content pane and drop it on another content pane. The two panes switch places. Move all the panes to a two-column layout, or re-arrange the three-column layout.



TIP: If you have changed the layout to fewer than three columns, and you want to re-create an additional column, drag a content pane to the top right corner of the page.

Minimize and expand panes to control which information is given priority. As priorities change, change which content panes are minimized and expanded.

Using the SAN Status Page




When you log into a management group, the Status Page shows the current information for that management group. After the initial log in, the Status Page refreshes frequently, but not immediately. For example, after creating a new volume in a cluster, it takes several seconds for the new volume to appear on the Status Page.

When reviewing information in the content panes, many items will display an arrow to the left of the item. Clicking on that arrow jumps to that item in the CMC, or, in a few cases, to the most closely related category. For example, clicking the arrow next to any of the Best Practice Summary items jumps to the management group containing those Best Practice items.

The SAN Status Page contains toolbar buttons to dock and undock the Status Page from the CMC, and to log into management groups. If the Status Page is undocked, and you close the CMC, when you open it next time, you must log into the management groups to retrieve the information for the Status Page. Use the button on the Status Page toolbar to log in to the management groups.

Alarms and events overview

The CMC displays the following types of events, based on their severity:

-  Informational—Provides status information about user activities (such as, creating a volume) and system actions and status (such as, a storage system is up). These types of events do not require taking action and are available only from the Events node for each management group.
-  Warning—Provides important information about a system component that may require taking action. These types of events are visible in both the Alarms window (for all management groups) and the Events node (for the management group where the alarm occurred).
-  Critical—Provides vital information about a system component that requires user action. These types of events are considered to be alarms and are visible in both the Alarms window (for all management groups) and the Events node (for the management group where the alarm occurred).

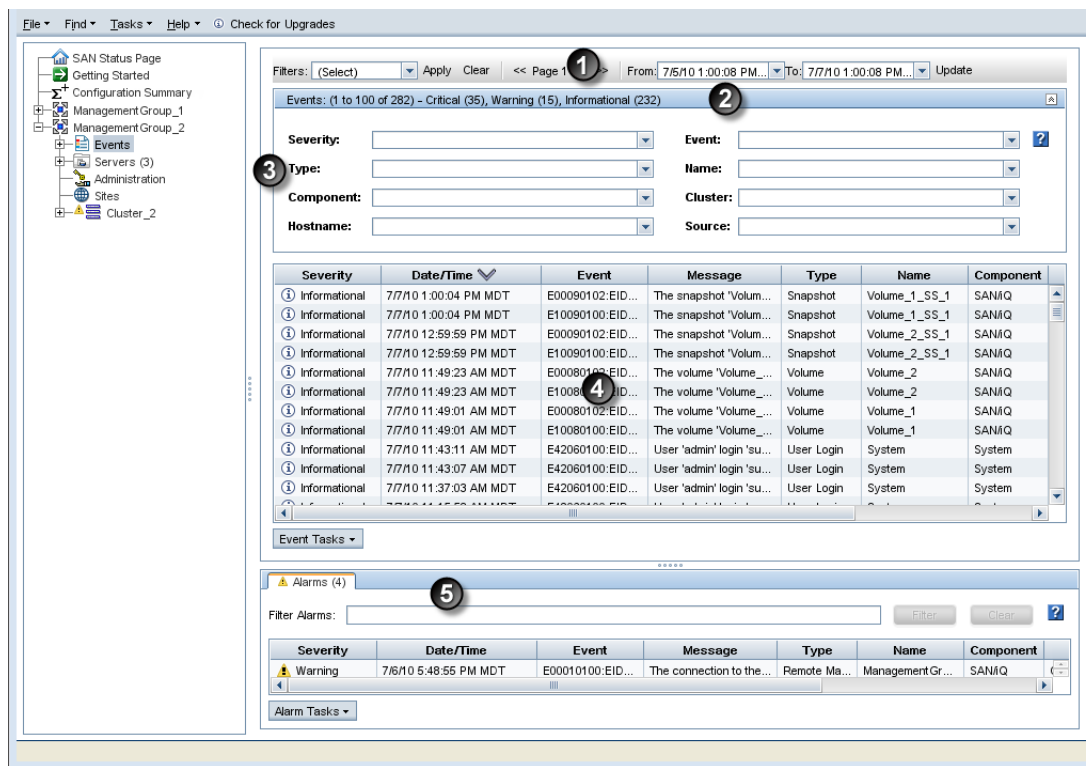
The system monitors over 400 different events and can send notification of events using email or SNMP or both.

NOTE: Be sure to set up notification of events for each management group. See [“Setting up email notification” \(page 91\)](#).

For more information about working with alarms and events, see [“Working with alarms” \(page 87\)](#) and [“Configuring events” \(page 88\)](#).

[Figure 39 \(page 86\)](#) shows the Events node and its parts.

Figure 39 Events node and its parts



1. Events toolbar
2. Number of events broken down by severity
3. Filter panel
4. Events list
5. Alarms list

NOTE: Except for the P4800 G2 SAN Solution for BladeSystem, alarms and events information is not available for storage systems listed under Available Systems in the CMC, because they are not currently in use on the SAN.

Table 28 (page 86) defines the alarms and events columns that appear in the CMC.

Table 28 Alarms and events column descriptions

Column	Description
Severity	Severity of the event: informational, warning, or critical.
Date/Time	Date and time the event occurred.
Event	Identification number and name of event.
Message	Brief text that describes the event.
Type	Specific type of device or object the event is related to, such as, volume, storage system, disk, manager, and so forth.
Name	Name of the device or object the event is related to.
Component	Major category of the device or object the event is related to, typically, hardware, SAN/iQ, or system (software).
Cluster	Name of the cluster where the event occurred. If blank, the cluster name was not known when the event occurred.
Management Group	Name of management group where the event occurred. If blank, the management group name was not known when the event occurred. (Only appears in the Alarms window.)

Table 28 Alarms and events column descriptions *(continued)*

Column	Description
Hostname	Name of the storage system that posted the event.
Source	User that generated the event, such as system (software) or a user name.

Working with alarms

The Alarms window is always visible at the bottom of the CMC and displays all warning and critical events for all management groups you are logged in to. Use the column headings to sort the list of alarms. Double-click an alarm to see more details.

Review alarms regularly, especially the critical ones, and take the appropriate action to resolve the alarm. Alarms stay in the Alarms window until the situation is resolved. A warning or critical icon also appears in the navigation window tree next to the item that has the warning or critical event.

For definitions of the event types and columns, see [“Alarms and events overview” \(page 85\)](#).

NOTE: An Alerts tab will also appear in two configurations: if storage systems in a management group are running a mix of SAN/iQ software versions, including pre-9.0 versions, or if monitoring multiple management groups, at least one of which is running pre-9.0 software. In these cases, only the Alerts tab will display alarms and events for the management group. When all storage systems are upgraded to 9.0 or later, the Alerts tab is not shown, and the Events node and Alarms tab will reflect all alarms and events.

Filtering the alarms list

1. In the navigation window, log in to the management group.
2. In the Alarms window, enter text to use as a filter in the Alarm Filter field.
The text must appear somewhere in one or more alarms and is case sensitive.
3. Click **Filter**.
The list of alarms changes to display only those that contain the filter text.
4. To display all alarms, click **Clear** to remove the filter.

Viewing and copying alarm details

1. In the navigation window, log in to the management group.
2. In the Alarms window, double-click an alarm.
3. For assistance with resolving the alarm, click the link in either the Event field or the Resolution field.
The link opens to a database that contains advisories and documents that may have additional information about the event and how to resolve it. If no results are found, no advisories that directly apply to that event have been published yet.
4. Click **Previous Alarm** or **Next Alarm** to view the details of other alarms.
5. To paste the event details into a document or email message, click **Copy** to copy the details to the clipboard.
6. Click **Close** to close the Alarm details window.

Viewing alarms in a separate window

View the alarms in a separate window that you can resize and move to a convenient location on the screen. Filter alarms, view alarm details, and export alarms from this window.

1. In the navigation window, log in to the management group.
2. In the Alarms window, click **Alarm Tasks**, and select **Open Alarms in Window**.
3. To close the window, click **Close**.

Exporting alarm data to a .csv file

1. In the navigation window, log in to the management group.
2. In the Alarms window, click **Alarm Tasks**, and select **Export Alarms**.
3. In the Filename field, enter a path and file name.
If you enter just a file name, the file is stored in the directory that the CMC is installed in.
4. Click **OK**.

Configuring events

The Events node displays all events for the current management group, according to the dates in the From and To fields in the toolbar. By default, the system displays about 1,000 events, broken down into pages of 100 events. Scroll through the current page of events and use the << and >> buttons to display different pages. The title bar of the Filters panel lists the number and severity of the events, shown in [Figure 39 \(page 86\)](#).

For definitions of the event types and columns, see [“Alarms and events overview” \(page 85\)](#).

Changing the event retention period

NOTE: Because of space constraints, retention periods longer than one month are not guaranteed.

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. Click **Event Tasks**, and select **Edit Event Log Policy**.
4. Change the event retention period. The default is one month.
5. Click **OK**.

Setting up remote log destinations

Use remote log destinations to automatically write all events for the management group to a computer other than the storage system. For example, direct the event data to a single log server in a remote location.

You must also configure the destination computer to receive the log files by configuring syslog on the destination computer. The syslog facility to use is local5, and the syslog levels are LOG_INFO, LOG_WARNING, LOG_CRIT. See the syslog documentation for that computer for information about configuring syslog.

To set up remote log destinations:

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. Click **Event Tasks**, and select **Edit Remote Log Destinations**.
4. In the Destinations field, enter the IP address or host name of the computer that will receive the events.
5. Select the event severities to include.
6. To set up a second remote destination, enter the appropriate information.
7. Click **OK**.

Viewing events in a separate window

View the events in a separate window. Resize and move the window to a convenient location on your screen.

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. Click **Event Tasks**, and select **Open Events in Window**.
The events open in a separate window. Use this window to filter, view details, and export events.
4. To close the window, click **Close**.

Working with events

The Events node offers multiple ways to manage and use event data. Filter events to control the information displayed, and export the data for use in a spreadsheet or other analysis tool.

Viewing new events

When new events occur after you log in to a management group, (New) appears next to the Events node.

To view new events:

1. Click **New Events** on the toolbar to bring in the newest events.
2. Sort by the Date/Time column to view the newest events.

Filtering the events list

To filter events you must first log in to the management group and select **Events** in the tree. Filter events listed in the Events node the following ways:

- Using the Filters list
- Changing the date range
- Using the filters panel

To use the Filters list:


1. From the Filters list, select an option to filter on.
Options in bold are predefined filters you cannot change. Options that are not bold are custom filters that you have saved from the filters panel, described in [“Saving filter views” \(page 90\)](#).
2. Click **Apply**.
To remove the filter, click **Reset**.

To change the date range:

1. In the From list, select **Choose From**, and select the date.
2. Click **OK**.
3. In the To list, select **Choose To**, and select the date.
4. Click **OK**.
5. Click **Update**.

Combine these date range filters with the options available in the filters panel described below.

To use the filters panel:

1. In the Events window, open the filters panel by clicking the expand button  (right side below the toolbar).
2. Use the filter lists to narrow the list of events.

If you select options from more than one filter list, the system does an “and,” so that all of the selected options must apply to an event for the event to stay in the Events list.

3. Click **Apply** on the Events toolbar.

The events list displays only those events matching the selected criteria.

To remove the filters and view all events, click **Reset**.

Saving filter views

Save filter views for later use. Custom filters appear in the Filters list in the toolbar below the boldface list of generic filters. Custom filters are available for use with any management group from this installation of the CMC.

1. Click **Events Tasks**, and select **Save Current View as Filter**.
2. Give the filter settings a name and click **Save**.

Deleting custom filters

1. Click **Event Tasks**, and select **Delete Filter**.
2. Select the custom filters to delete.
3. Click **OK**.
4. Click **Delete** to delete the selected custom filters.

Viewing event details

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. In the Events list, double-click an event.
4. For assistance with resolving the event, click the link in either the Event field or the Resolution field.

The link opens a database that contains advisories and documents that may have additional information about the event and how to resolve it. If no results are found, no advisories that directly apply to that event have been published yet.

5. Click **Previous Event** or **Next Event** to display the details of other events.
6. (Optional) To paste the event details into a document or email message, click **Copy** to copy the details to the clipboard.
7. Click **Close** to close the Event Details window.

Copying events to the clipboard

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. Do one of the following:
 - Select one or more events, click **Event Tasks**, and select **Copy Selected to Clipboard**.
 - Click **Event Tasks**, and select **Copy All to Clipboard**.

Exporting event data to a .csv or .txt file

1. In the navigation window, log in to the management group.
2. Select **Events** in the tree.
3. Click **Event Tasks**, and select **Export Events**.
4. In the Filename field, enter a path and file name.

If you enter just a file name, the file is stored in the directory that the CMC is installed in. To select a different location or .txt format, click **Browse**.

5. Click **OK**.

Setting up email notification

To set up email notification of events, you must set up the following:

- Email server
- Email recipients

Setting up the email server

When you configure a management group, you configure email notification for events by setting up the email server to send events and then adding email recipients to receive notification based on event severity. Change these settings later, if necessary, by using the Email window found under the Events node in the management group tree.

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**Email**.
3. Click **Email Tasks**, and select **Edit Email Settings**.
4. In the Server IP/Hostname field, enter the IP address or host name of your email (SMTP) server.
5. In the Server Port field, enter the email port.

The standard port is 25.

6. In the Sender Address field, enter the email address, including the domain name, to use as the sender for notifications.

The system automatically adds the host name of the storage system in the email From field, which appears in many email systems. This host name helps identify where the event occurred.

7. Do one of the following:
 - To save your changes and close the window, click **Apply**.
 - To save your changes, close the window, and send a test email message, click **Apply and Test**.

If you do not complete the email configuration, an alarm appears until you do so. Your changes are saved, but the test function does not work until you complete the configuration, including at least one email recipient.

NOTE: In some configurations designed to prevent spam, mail servers may start dropping multiple emails if too many are being sent from a single source. During some SAN/iQ operations, multiple alarms may be generated while the operation progresses. If you have a mail server configured to filter spam, you may not receive some emails.



TIP: Add the sender email address to your safe senders list to ensure the system emails are received.

Setting up email recipients

Set up email notification for events for each management group. After setting up the email server, add email recipients to receive notification based on event severity.

NOTE: To temporarily disable email notification to a recipient, possibly because you are doing something that would cause many event emails, you can deselect all of the severities. This generates an alarm that persists until you set at least one severity again for the recipient.

To set up email recipients:

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**Email**.
3. Click **Email Tasks**, and select **Edit Email Settings**.

4. Click **Add Recipient**.
5. In the Address field, enter the email address.
6. Select the following:
 - Severity of the events
 - Message type
 - Language
7. Click **OK**.
8. Click **Apply** or **Apply and Test**.

Setting up SNMP

The management group can be monitored using an SNMP client. You can also enable SNMP traps to receive system alerts. The Management Information Base (MIB) is read-only and supports SNMP versions 1 and 2c. See [“Installing the LeftHand Networks MIB” \(page 95\)](#) for a list of MIBs.

After installing SAN/iQ version 9.0, the SNMP agent on the storage system is enabled by default and allows read-only access using the community string `public`. You can change this configuration, unless you are using HP Insight Remote Support, which requires the community string must be set to `public`. To receive notification of events, you must configure SNMP traps.

If using the HP System Management Homepage, view the SNMP settings there. You can also start SNMP and send test v1 traps.

Enabling SNMP agents

Most storage systems allow enabling and disabling SNMP agents. After installing version 9.0, SNMP will be enabled on the storage system by default.

Configuring SNMP includes these tasks:

- Enabling the SNMP Agent and adding a community string, if necessary
The community string acts as an authentication password. It identifies hosts that are allowed read-only access to the SNMP data. The community `public` typically denotes a read-only community. This string is entered into an SNMP client when attempting to access the system.
- Configuring access control

Enabling the SNMP agent

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**SNMP**.
3. Click **SNMP Tasks** and select **Edit SNMP General Settings**.
4. Select the **Enabled** radio button to activate the SNMP Agent fields.
5. Enter the Community String. If using Insight Remote Support, the community string must be set to Public.
6. [Optional] Enter System Location information for the storage system.
For example, this information may include the address, building name, room number, and so on.
7. [Optional] Enter System Contact information.
Normally this will be the SAN/iQ administrator information, such as email address or phone number for the person to contact about the storage system.
8. Click **OK**.
9. Continue with configuring access control for SNMP clients.

Configuring access control for SNMP clients

Enable the SNMP agent and configure access control for SNMP clients. Enter either a specific IP address and the IP Netmask as None to allow a specific host to access SNMP, or specify the Network Address with its netmask value so that all hosts matching that IP and netmask combination can access SNMP.

NOTE: Use the CMC ping feature to verify IP addresses while configuring access control. See “Pinging an IP address” (page 54).

Adding an SNMP client

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**SNMP**.
3. Click **SNMP Tasks**, and select **Edit SNMP General Settings**.
4. In the Access Control section, click **Add** to add an SNMP client that to use for viewing SNMP. Add SNMP clients by specifying either IP addresses or host names. For HP remote support, add the CMS for HP Insight Remote Support.
5. Do one of the following:
 - Select **By Address** and enter the IP address, then select an IP Netmask from the list. Select Single Host if adding only one SNMP client.
 - Select **By Name** and enter a host name.
That host name must exist in DNS and the management group must be configured with DNS for the client to be recognized by the host name.
6. Click **OK** to add the entry to the Access Control list.
7. Click **OK** in the Edit SNMP Settings window to finish.

Editing access control entries

After installing version 8.5 or later, by default, access control is open from any system using the “public” community string. This access control entry is listed as “default.” If you delete this entry and want to add it back in, use the By Name option and enter “default” as the name.

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**SNMP**.
3. Click **SNMP Tasks**, and select **Edit SNMP General Settings**.
4. Select the Access Control entry from the list.
5. Click **Edit**.
6. Change the appropriate information.
7. Click **OK**.
8. Click **OK** on the Edit SNMP Settings window when finished.

Deleting access control entries

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**SNMP**.
3. Click **SNMP Tasks** and select **Edit SNMP General Settings**.
4. Select a client listed in the Access Control list and click **Delete**.
5. Click **OK** to confirm.
6. Click **OK** on the Edit SNMP Settings window when finished.

Disabling the SNMP agent

Disable the SNMP Agent if no longer using SNMP applications to monitor the network of storage systems.

Disabling SNMP

1. In the navigation window, log in to the management group.
2. In the tree, select **Events→SNMP**.
3. Click **SNMP Tasks**, and select **Edit SNMP General Settings**.
4. Select **Disable SNMP Agent**.

Note that the Agent Status field now shows disabled. The SNMP client information remains listed, but cannot be used.

Adding SNMP traps

Verify that SNMP is enabled, which it is by default.

Add a Trap Community String, which is used for client-side authentication, and trap recipients.

1. In the navigation window, log in to the management group.
2. In the tree, select **Events→SNMP**.
3. Click **SNMP Tasks** and select **Edit SNMP Traps Settings**.
4. Enter the Trap Community String.

The trap community string does not have to be the same as the community string used for access control, but it can be.

5. Click **Add** to add trap recipients.
6. Enter the IP address or host name for the SNMP client that is receiving the traps.
For HP remote support, add the CMS for HP Insight Remote Support.
7. Select the Trap Version.
Version 1 is required for HP remote support.
8. Click **OK**.
9. Repeat steps 5 through 8 for each trap recipient.
10. Click **OK** on the Edit SNMP Traps window when finished adding hosts.

Editing trap recipients

1. Log in to the storage system and expand the tree.
2. In the navigation window, log in to the management group.
3. In the tree, select **Events→SNMP**.
4. Click **SNMP Tasks**, and select **Edit SNMP Traps Settings**.
5. Select one of the Trap Recipients, and click **Edit**.
6. Change the information as needed.
Trap Version 1 is required for HP remote support.
7. Click **OK**.
8. Click **OK** when finished editing trap recipients.

Removing trap recipients

1. In the navigation window, log in to the management group.
2. In the tree, select **Events→SNMP**.
3. Click **SNMP Tasks**, and select **Edit SNMP Traps Settings**.
4. Select one of the Trap Recipients, and click **Remove**.
5. Click **OK** on the SNMP Traps tab when finished removing trap recipients.

Sending a test trap

Send a test trap to verify that your trap recipients are working.

1. In the navigation window, log in to the management group.
2. In the tree, select **Events**→**SNMP**.
3. Click **SNMP Tasks**, and select **Send Test Trap**.
4. Click **OK** on the Test SNMP Traps message window.

Disabling SNMP traps

To disable SNMP traps, you must delete all of the settings in the SNMP Traps window.

1. Remove the Trap Recipient hosts.
2. Delete the Trap Community String.
3. Click **OK**.

Using the SNMP MIBs

The LeftHand Networks MIBs provide read-only access to the storage system. The SNMP implementation in the storage system supports MIB-II compliant objects.

These files, when loaded in the SNMP client, allow you to see storage system-specific information such as model number, serial number, hard disk capacity, network characteristics, RAID configuration, DNS server configuration details, and more.

NOTE: With version 8.5 and later, traps no longer all use a single OID. The LEFTHAND-NETWORKS-NOTIFICATION-MIB defines the OIDs now in use.

Installing the LeftHand Networks MIB

The complete set of standard SNMP MIB files and the LeftHand Networks MIB files are installed when installing the CMC using the Complete option. The installer places the MIBs in the following directory by default: C:\Program Files\HP\P4000\UI\mibs. Your SNMP client may require that you copy the MIBs to another location, or you may need to copy them to the system where your SNMP client is installed.

On the system where your SNMP client is installed, load the LeftHand Networks MIBs as outlined below, using the SNMP client. The complete set of standard SNMP MIBs must also be loaded.

Load the MIBs as follows:

1. If you do not have the standard SNMP MIBs loaded, load them.
2. If you do not have HCNUM-TC.MIB loaded, load it.
3. Load LEFTHAND-NETWORKS-GLOBAL-REG-MIB.
4. Load LEFTHAND-NETWORKS-NSM-MIB.
5. The following MIB files can be loaded in any sequence.
 - LEFTHAND-NETWORKS-NSM-CLUSTERING-MIB
 - LEFTHAND-NETWORKS-NSM-DNS-MIB
 - LEFTHAND-NETWORKS-NSM-INFO-MIB
 - LEFTHAND-NETWORKS-NSM-NETWORK-MIB
 - LEFTHAND-NETWORKS-NSM-NOTIFICATION-MIB
 - LEFTHAND-NETWORKS-NSM-NTP-MIB
 - LEFTHAND-NETWORKS-NSM-SECURITY-MIB
 - LEFTHAND-NETWORKS-NSM-STATUS-MIB
 - LEFTHAND-NETWORKS-NSM-STORAGE-MIB

The supported MIBs

The following are the supported standard MIBs, though not every function in each MIB is supported.

- DISMAN-EVENT-MIB
- HOST-RESOURCES-MIB
- IF-MIB
- IP-FORWARD-MIB
- IP-MIB
- NET-SNMP-AGENT-MIB
- NET-SNMP-EXTEND-MIB
- NETWORK-SERVICES-MIB
- NOTIFICATION-LOG-MIB
- RFC1213-MIB
- SNMP-TARGET-MIB
- SNMP-VIEW-BASED-ACM-MIB
- SNMPv2-MIB
- UCD-DLMOD-MIB
- UCD-SNMP-MIB

Running diagnostic reports

Use diagnostics to check the health of the storage system hardware. Different storage systems offer different sets of diagnostic tests.

NOTE: Running diagnostics helps monitor the health of the storage system or to troubleshoot hardware problems.

1. Select a storage system in the navigation window.
2. Open the tree below the storage system and select **Diagnostics**.
3. On the Diagnostics tab, review the list of diagnostic tests available.

The default setting is to run all tests. Customize the set of tests to run by changing which tests are selected. Clear the check box next to a test to stop that test from running. Click **Diagnostics Tasks** and select **Check All** or **Clear All** to streamline the selection process.

NOTE: Running all of the diagnostic tests will take several minutes. To shorten the time required to run tests, clear the check boxes for any unneeded tests.

4. Click **Diagnostic Tasks** and select **Run Tests**.
A progress message appears. When the tests complete, the results of each test appear in the Result column.
5. (Optional) When the tests complete, to view a report of test results, click **Diagnostic Tasks** and select **Save to File**.
6. Select a location for the diagnostic report file and click **Save**.

NOTE: If any of the diagnostics show a result of “Failed,” call Customer Support.

List of diagnostic tests

Table 29 (page 97) shows a sample of the diagnostic tests that are available for the storage system, including the description of that test and the pass/fail criteria.

For each test, lists the following information:

- A description of the test
- Pass / fail criteria

NOTE: Available diagnostic tests depend on the storage system.

For VSA, only the Disk Status Test is available.

Table 29 Example list of hardware diagnostic tests and pass/fail criteria

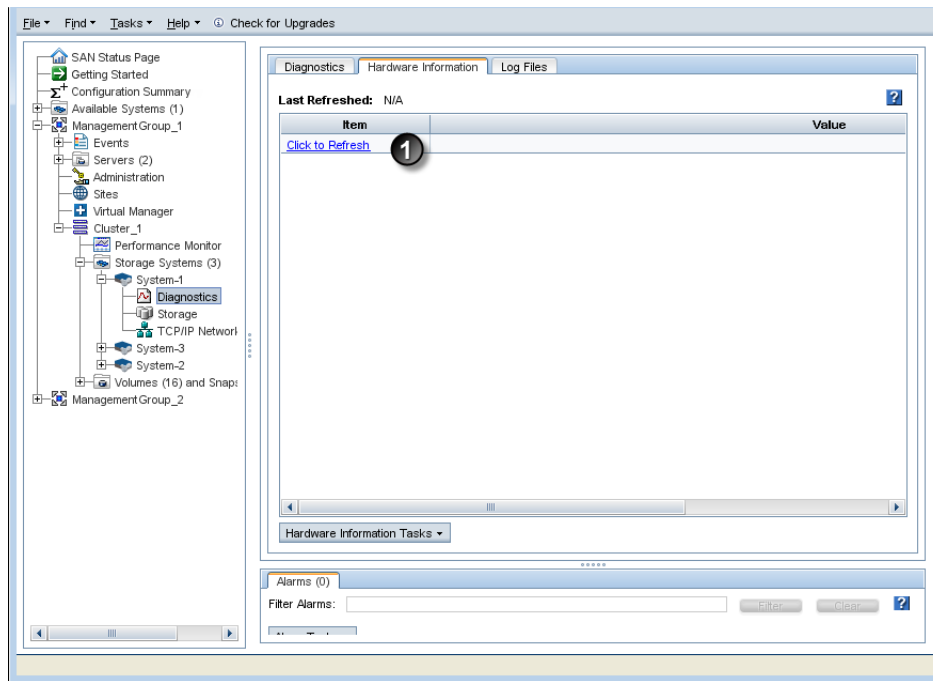
Diagnostic test	Description	Pass criteria	Fail criteria
Fan Test	Checks the status of all fans.	Fan is normal	Fan is faulty or missing
Power Test	Checks the status of all power supplies.	Supply is normal	Supply is faulty or missing
Temperature Test	Checks the status of all temperature sensors.	Temperature is within normal operating range	Temperature is outside normal operating range
Cache Status	Checks the status of the disk controller caches.	Cache is normal	Cache is corrupt
Cache BBU Status	Checks the status of the battery backed-up cache.	The BBU is normal and not charging or testing	The BBU is charging, testing or faulty
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing
Disk Temperature Test	Checks the temperature of all disk drives.	The temperature is within normal operating range	The temperature is outside normal operating range
Disk SMART Health Test	S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks.	All drives pass health test	Warning or Failed if one or more drives fails health test
Generate SMART logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated
Generate DSET Report & Perc Event Logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated
Generate Platform or HP Diagnostic Report (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated
Generate IBM Support logs (for analysis contact IBM Support)	Generates IBM Support logs when requested by Customer Support.	The logs were successfully generated	The logs were not generated

Generating a hardware information report

Hardware information reports display statistics about the performance of the storage system, its drives and configuration. Statistics in the hardware reports are point-in-time data, gathered by clicking the Refresh button on the Hardware Information tab.

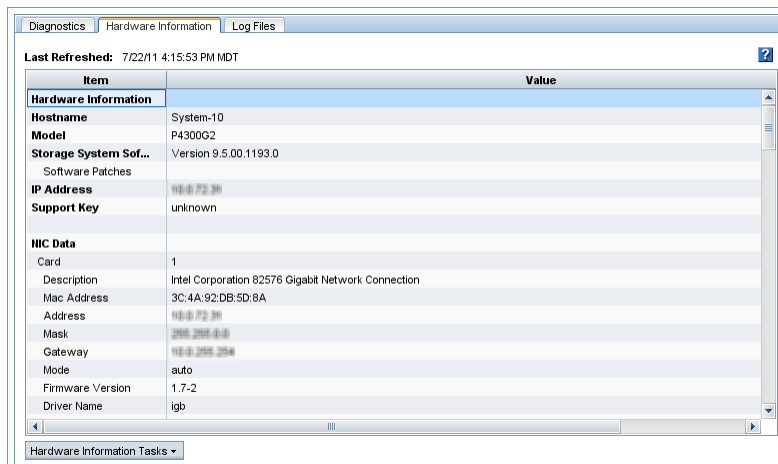
1. Select the **Hardware Information** tab.

Figure 40 Opening the hardware information window



1. Link to obtain hardware statistics
2. On the Hardware table, use the link **Click to Refresh** to obtain the latest hardware statistics.

Figure 41 Viewing the hardware information for a storage system



Saving a hardware information report

1. Click **Hardware Information Tasks** and select **Save to File** to download a text file of the reported statistics.
2. Choose the location and name for the report.
3. Click **Save**.

The report is saved with an .html extension.

Hardware information report details

Available hardware report statistics vary depending on the storage system.

Table 30 Selected details of the hardware report

This term	means this
Hostname	Hostname of the storage system.
Storage system software	Full version number for storage system software. Also lists any patches that have been applied to the storage system.
IP address	IP address of the storage system.
Support key	Support Key is used by a Technical Support representative to log in to the storage system.
NIC data	Information about NICs in the storage system, including: <ul style="list-style-type: none">• Card number• Manufacturer description• MAC address• IP address• Mask• Gateway• Mode• Firmware version• Driver name• Driver version
DNS data	Information about DNS, if a DNS server is being used, providing the IP address of the DNS servers. IP address of the DNS servers.
Memory	Information about RAM in the storage system, including values for total memory and free memory in GB.
CPU	Details about the CPU, including model name or manufacturer of the CPU, clock speed of the CPU, and cache size.
Stat	Information about the CPU. CPU seconds shows the number of CPU seconds spent on user tasks, kernel tasks, and in idle state. Machine uptime is the total time the storage system has been running from initial boot up.
Backplane information	Selected information about the backplane LEDs: LED support and id LED.
Motherboard information	Includes chassis serial number and BIOS version
Drive info	For each drive, reports the model, serial number, and capacity.
Drive status	For each drive, reports the status, health, and temperature.
RAID	Information about RAID. This includes: <ul style="list-style-type: none">• Rebuild rate• Unused devices• Statistics• Unit number
RAID O/S partitions	Information about O/S RAID.
Boot-device statistics	Status information about the boot device: status, capacity in MB, driver version, media used for device, and model.
Statistics	Information about the O/S RAID for the storage system.

Table 30 Selected details of the hardware report *(continued)*

This term	means this
Unit number	Identifies devices that make up the O/S RAID configuration, including: <ul style="list-style-type: none">• Type of storage (BOOT, LOG, SANIQ, DATA)• RAID level (0, 1, 5)• Status (Normal, Rebuilding, Degraded, Off)• Capacity• Rebuild statistics (% complete, time remaining)
Controller/cache items	Depending on the storage system, this can include: <ul style="list-style-type: none">• Information about the RAID controller card and Battery Backup Unit (BBU)• Information about RAM.
Power supply	Shows the type or number of power supplies.
Power supplies	Status information about the power supplies.
Sensors	Shows for the hardware listed, the status, real measured value, minimum and maximum values.

Using log files

If Technical Support requests that you send a copy of a log file, use the Log Files tab to save that log file as a text file.

The Log Files tab lists two types of logs:

- Log files that are stored locally on the storage system (displayed on the left side of the tab).
- Log files that are written to a remote log server (displayed on the right side of the tab). This list is empty until you configure remote log files and the remote log target computer.

Saving log files locally

1. Select a storage system in the navigation window.
2. Open the tree below the storage system and select **Diagnostics**.
3. Select the **Log Files** tab.
4. To retrieve the latest data, click **Log File Tasks** and select **Refresh Log File List**.
5. Scroll down the list of log files in the Log Files list and select the file or files to save.
To select multiple files, press **Ctrl** or **Ctrl+Shift**.
6. Click **Log Files Tasks** and select **Save Log Files**.
7. Select a location for the file or files.
8. Click **Save**.

Exporting the System Summary

The System Summary has information about all of the storage systems on the network. Export the summary to a .csv file for use in a spreadsheet or database.

Information in the summary includes:

- Storage system information, [“Working with storage systems”](#) (page 21)
- Management group information, [“Working with management groups”](#) (page 103)
- Network information, [“Managing the network”](#) (page 48)
- RAID information, [“Storage Configuration: Disk RAID and Disk Management”](#) (page 30)

To export the summary:

1. From the CMC menu bar, select **Tasks→System Summary**.
2. Click **Export**.
3. Select a location for the file, and rename it if desired.
4. Click **Export**.

Configuring a remote log and remote log destination

Use remote log files to automatically write log files to a computer other than the storage system. For example, direct the log files for one or more storage systems to a single log server in a remote location. The computer that receives the log files is called the Remote Log Target.

You must also configure the target computer to receive the log files.

1. Select a storage system in the navigation window.
2. Open the tree below the storage system and select **Diagnostics**.
3. Select the **Log Files** tab.
4. Click **Log File Tasks** and select **Add Remote Log Destination**.
5. In the Log Type list, select the log to direct to a remote computer.

The Log Type list only contains logs that support syslog.

6. In the Destination field, enter the IP address or host name of the computer that will receive the logs.

For a Windows operating system, find out the name of the remote computer with Control Panel > System Properties > Computer Name.

7. Click **OK**.

The remote log appears in the Remote logs list on the Log Files window.

Configuring the remote log target computer

Configure syslog on the remote log target computer. Refer to the syslog product documentation for information about configuring syslog.

NOTE: The string in parentheses next to the remote log name on the Log Files tab includes the facility and level information that you will configure in syslog. For example, in the log file name: auth error (auth.warning) the facility is "auth" and the level is "warning."

Editing remote log targets

Select a different log file or change the target computer for a remote log:

1. Select a storage system in the navigation window.
2. Open the tree below the storage system and select **Diagnostics**.
3. Select the **Log Files** tab.
4. Select the log in the Remote logs list.
5. Click **Log File Tasks** and select **Edit Remote Log Destination**.
6. Change the log type or destination and click **OK**.
7. Ensure that the remote computer has the proper syslog configuration.

Deleting remote logs

1. Select a storage system in the navigation window.
2. Open the tree below the storage system and select **Diagnostics**.
3. Select the **Log Files** tab.
4. Click **Log File Tasks** and select **Delete Remote Log Destination**.
5. Click **OK** on the confirmation window.

NOTE: After deleting a remote log file from the storage system, remove references to this log file from the syslog configuration on the target computer.

Exporting support logs

If asked to do so by Customer Support, export support logs for a management group or storage system. The logs are saved as a zip file.

1. Select a management group or storage system in the navigation window.
2. Do one of the following depending on what you selected:
 - Click **Management Group Tasks** and select **Export Management Group Support Bundle**.
 - Click **Storage System Tasks** and select **Export Storage System Support Bundle**.
3. Select the location to save the file.

You cannot change the name of the zip file.
4. Click **Save**.

8 Working with management groups

A management group is a collection of one or more storage systems. It is the container within which you cluster storage systems and create volumes for storage. Creating a management group is the first step in creating an IP SAN with the SAN/iQ software.

Functions of management groups

Management groups serve several purposes:

- **Management groups are the highest administrative domain for the SAN.** Typically, storage administrators will configure at least one management group within their data center.
- **Organize storage systems into different groups for categories of applications and data.** For example, you might create a management group for Oracle applications and a separate management group for Exchange.
- **Ensure added administrative security.** For example, you could give the system administrator in charge of Exchange access to the Exchange management group but not the Oracle management group.
- **Prevent some storage resources from being used unintentionally.** If a storage system is not in a management group, the management group cannot use that storage system as a storage resource. For example, all of the storage systems in a management group can be pooled into clusters for use by volumes in that group. To prevent a new storage system from being included in this pool of storage, put it in a separate management group.
- **Contain clustering managers.** Within a management group, one or more of the storage systems acts as the managers that control data transfer and replication.

Guide for management groups

When using the Management Groups, Clusters and Volumes wizard, you must configure the characteristics described in [Table 31 \(page 103\)](#). When the management group is created, check the Best Practice Summary to verify that the configuration is following best practices for availability and data protection. See [“Best Practice summary overview” \(page 111\)](#).

Table 31 Management group requirements

Management group requirement	What it means
Configure storage systems	Before you create a management group, make sure you know the IP addresses for the storage systems for the cluster. Also, make sure they are configured for network bonding as best fits your network environment. CAUTION: [VSA] You cannot clone a VSA after it is in a management group. You must clone a VSA while it is in the Available Systems pool.
Plan administrative users	When you create a management group, you must add the first administrative user. This user has full administrative permissions. Add additional users after the management group is created. See “Adding a new administrative user” (page 79) .
Plan date and time configuration	You can use an NTP server or manually set the date, time, and time zone for the management group. You should know the configuration you want to use before beginning the wizard. See “Setting the date and time” (page 76) .
Plan DNS configuration	You can configure DNS at the management group level for all storage systems in the management group. The storage system can use a DNS server to resolve host names.

Table 31 Management group requirements *(continued)*

Management group requirement	What it means
	You need the DNS domain name, suffix, and server IP address. See “Using a DNS server” (page 71).
Plan email notification	You can set up email notification for events for each management group. You must set up the email server to send events. You need the email (SMTP) server IP or host name, server port, a valid email address to use as the sender address. See “Setting up email notification” (page 91).
Plan type of cluster	A cluster can be standard or Multi-Site. If you want to create a Multi-Site configuration, you need the physical sites and the storage systems that go in them already created.
Plan virtual IP addresses (VIPs)	Plan a unique VIP for each cluster. VIPs ensure fault-tolerant server access to the cluster and enable iSCSI load balancing. You need the IP address and subnet mask. See “Virtual IP addresses” (page 229).
[Optional] Plan volume size and data protection	If you create a volume in the wizard, you need the volume size and data protection level. See “Guide for volumes” (page 154).

Creating a management group

Creating a management group is the first step in the process of creating clusters and volumes for storage. Tasks included in creating a management group are:

- Planning the management group configuration
- Creating the management group by using the Management Groups, Clusters and Volumes wizard
- Ensuring you have the proper configuration of managers

Creating a new management group

1. Select Getting Started in the navigation window to access the Getting Started Launch Pad.
2. Select the **Management Groups, Clusters and Volumes Wizard**.
3. [Optional] Click the link to review the information you will need to have ready to create the management group and cluster.
4. Click **Next** to start creating the management group.
5. Select **New** to create a new management group and click **Next**.

Create management group and add storage systems

1. Enter a name for the new management group.
This name cannot be changed later without destroying the management group.
2. Select the storage system(s) to add to the management group.
Use **Ctrl+Click** to select more than one.

- ① **IMPORTANT:** Using two storage systems in a management group requires a Failover Manager for the highest level of data availability. If you are creating a management group with two storage systems and you do not add a FOM, you must select the checkbox acknowledging the risk before you can continue.

Add administrative user

1. Click **Next** to add an administrative user.
2. Enter the administrative user's name, a description, and a password.
The first administrator is always at full administrator level.
3. Click **Next** to set the time for the management group.

Set management group time

1. Select the method by which to set the management group time.
 - [Recommended] To use an NTP server, know the URL of the server, or its IP address, before you begin.
Note: if using a URL, DNS must be configured on the storage systems in the group.
 - To set the time manually, select **Edit** to display the Date and Time Configuration window. Check each field on this window to set the time for all storage systems in this management group.
2. Click **Next** to set the DNS server.

Set DNS server

1. Enter the DNS domain, suffix, and server IP address.
2. Click **Next** to set up email for notification.

Set up email for notification

1. Enter the email (SMTP) server IP address or host name, port number, and email address to use for the sender for event notification.
2. Click **Next** to create a cluster.

Create cluster and assign a VIP

The following steps are for creating a standard cluster. If you are creating a Multi-Site cluster, see "Creating Multi-Site Clusters and Volumes" in Chapter 2 of the *HP P4000 Multi-Site HA/DR Solution Pack User Guide*.

1. Select **Standard Cluster** on the Create a Cluster window, and click **Next**.
2. Enter a cluster name in the Create Cluster window.
3. From the list, select the storage systems to include in the cluster.
4. Click **Next** to assign a Virtual IP.
5. Add the VIP and subnet mask.
6. Click **Next** to create a volume and finish creating the management group.

Create a volume and finish creating management group

1. Enter a name, description, data protection level, size, and provisioning type for the volume.
2. Click **Finish**.
3. Review the details on the summary window and click **Close**.
4. A message opens notifying you to register and receive license keys.
5. Click **OK**.
6. As a last step, save a .txt file of the configuration information for the entire management group.
See ["Saving management group configuration information"](#) (page 116).

Management group map view tab

After you create the management group and finish setting up the SAN, use the Map View tab for viewing the relationships between servers, sites, clusters, volumes, snapshots, and remote copies

in the management group. For more information on using the map view tools, see [“Using the Map View ” \(page 14\)](#).

Best practice for managers in a management group

When creating a management group, the wizard creates an optimal manager configuration for the number of storage systems used to create the group. See [Table 32 \(page 106\)](#) for the default manager configurations.

After you have finished creating the management group, be certain to reconfigure managers as necessary to optimize your particular SAN configuration. The Best Practice analyzer for the management group indicates if the manager configuration is satisfactory. For more information about managers, see [“Managers overview” \(page 106\)](#).

Table 32 Default number of managers added when a management group is created

Number of storage systems	Manager configuration
1	1 manager
2	2 managers and a Virtual Manager
3 or more	3 managers

Managers overview

Within a management group, managers are storage systems that govern the activity of all of the storage systems in the group. All storage systems contain the management software, but you must designate which storage systems run that software by starting managers on them. These storage systems then “run” managers, much as a PC runs various services.

Functions of managers

Managers have the following functions:

- Control data replication. (Note: managers are not directly in the data path.)
- Manage communication between storage systems in the cluster.
- Resynchronize data when storage systems change states.
- Coordinate reconfigurations as storage systems are brought up and taken off line.

One storage system has the coordinating manager. You can determine which storage system is the coordinating manager by selecting the management group, then clicking the Details tab. The Status field at the top shows the coordinating manager.

Managers and quorum

Managers use a voting algorithm to coordinate storage system behavior. In this voting algorithm, a strict majority of managers (a quorum) must be running and communicating with each other in order for the SAN/iQ software to function. An odd number of managers is recommended to ensure that a majority is easily maintained. An even number of managers can get into a state where no majority exists—one-half of the managers do not agree with the other one-half. This state, known as a “split-brain,” may cause the management group to become unavailable.

For optimal fault tolerance in a single-site configuration, you should have 3 or 5 managers in your management group to provide the best balance between fault tolerance and performance. The maximum supported number of managers is 5. See [Table 33 \(page 107\)](#).

Table 33 Managers and quorum

Number of Managers	Number for a quorum	Fault tolerance	Explanation
1	1	None	If the manager fails, no data control takes place. This arrangement is not recommended.
2	2	None	Even number of managers not recommended, except in specific configurations. Contact Customer Support for more information.
3	2	High	If one manager fails, 2 remain, so there is still a quorum. (Note: 2 managers are not fault tolerant. See above.)
4	3	High	Even number of managers not recommended, except in specific configurations. Contact Customer Support for more information.
5	3	High	If one or two managers fail, 3 remain so there is still a quorum.

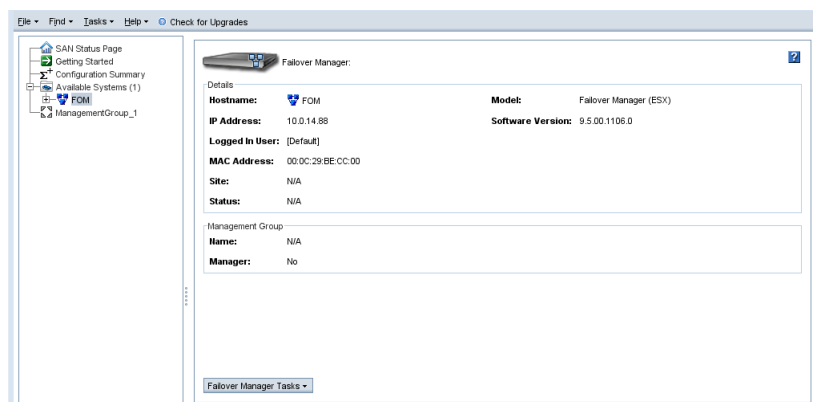
Regular managers and specialized managers

Regular managers run on storage systems in a management group. The SAN/iQ software has two other types of specialized managers, Failover Managers and Virtual Managers, described below. For detailed information about specialized managers and the how to use them, see [“Using specialized managers”](#) (page 120).

Failover Managers

The Failover Manager is used in two-system and in Multi-Site SAN configurations to support automatic quorum management in the SAN. Configuring a Failover Manager in the management group enables the SAN to have automated failover without requiring a regular manager running on a storage system. A Failover Manager runs as a virtual machine on a VMware Server or on ESX and must be installed on network hardware other than the storage systems in the SAN.

[Figure 42](#) (page 107) shows the Failover Manager installed, configured, and appearing in the CMC.

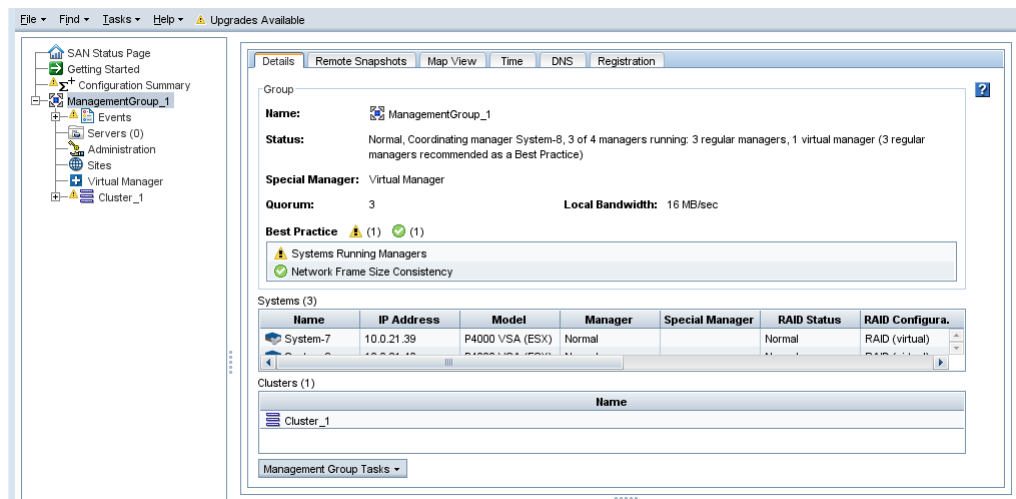
Figure 42 Failover Manager in the available systems pool

Once installed and configured, the Failover Manager operates as a storage system in how you add it to a management group where it serves solely as a quorum tie-breaking manager.

Virtual Managers

A Virtual Manager is added to a management group, as shown in [Figure 43 \(page 108\)](#), but is not started on a storage system until a failure in the system causes a loss of quorum. Unlike the Failover Manager, which is always running, the Virtual Manager must be started manually on a storage system after quorum is lost. It is designed to be used in two-system or two-site system configurations which are at risk for a loss of quorum.

Figure 43 Virtual Manager added to a management group



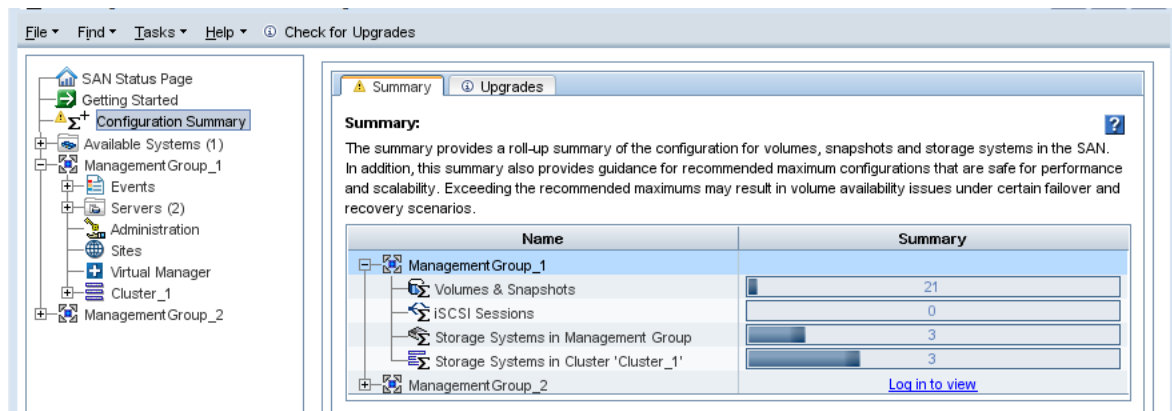
Configuration Summary overview

The Configuration Summary provides an easy-to-use reference for managing the size and optimum configuration of your SAN. The first time you create a management group, a configuration summary table is created that resides immediately below the Getting Started Launch Pad in the navigation window. Subsequent management groups are added to this Configuration Summary, shown in [Figure 44 \(page 109\)](#). For each management group, the Configuration Summary displays an overview of the volumes, snapshots, and storage systems in that management group. The Summary roll-ups display configuration information and guide you to optimal configurations for volumes and snapshots, iSCSI sessions, and the number of storage systems in the management group and in each cluster.

Summary roll-up

The summary roll-up provided on the Configuration Summary panel is organized by management group. Within each management group is listed the total number of volumes and snapshots, storage systems, and iSCSI sessions contained in the management group.

Figure 44 Configuration Summary created when the first management group is configured



Configuration guidance

As the Configuration Summary reports the numbers of the storage items, it provides warnings about the safe limits for each category, based on performance and scalability. These warnings first alert you that the category is nearing the limits by turning the category orange. When an individual category turns orange, the Configuration Summary category in the navigation window turns orange as well. When an individual category reaches the maximum recommended configuration it turns red. When the number in that category is reduced, the color changes immediately to reflect the new state. For example, if you have a large number of volumes that have numerous schedules that are creating and deleting snapshots, the snapshots may increase to a number that changes the summary bar from green to orange. As soon as enough snapshots from the schedules are deleted, reducing the overall total, the summary bar returns to green.

Best practices

The optimal and recommended number of storage items in a management group depend largely on the network environment, the configuration of the SAN, the applications accessing the volumes, and what you are using snapshots for. However, we can provide some broad guidelines that help you manage your SAN to obtain the best and safest performance and scalability for your circumstances. These guidelines are in line with our tested limits for common SAN configurations and uses. Exceeding these guidelines does not necessarily cause any problems. However, your performance may not be optimal, or in some failover and recovery situations may cause issues with volume availability.

Volumes and snapshots

The optimum number of combined volumes and snapshots ranges up to 1,000. If the management group contains 1,001 to 1,500 volumes and snapshots, the Configuration Summary appears orange for that line of the management group. Exceeding 1,500 volumes and snapshots triggers a warning by turning that line red. As soon as the total number reduces below the boundary, the summary bar returns to the previous indicator, either orange or green.

iSCSI sessions

The optimum number of iSCSI sessions connected to volumes in a management group ranges up to 4,000. If the management group contains 4,001 to 5,000 iSCSI sessions, the Configuration Summary appears orange for that line of the management group. Exceeding 5,001 iSCSI sessions triggers a warning by turning that line red. As soon as the total number of iSCSI sessions reduces below the boundary, the summary bar returns to the previous indicator, either orange or green.

Storage systems in the management group

The optimum number of storage systems in a management group ranges up to 20. If the management group contains 21 to 30 storage systems, the Configuration Summary appears orange for that line

of the management group. Exceeding 30 storage systems triggers a warning by turning that line red. As soon as the total number of storage systems reduces below the boundary, the summary bar returns to the previous indicator, either orange or green.

Storage systems in the cluster

The optimum number of storage systems in a cluster ranges up to 10. If the cluster contains 11 to 16 storage systems, the Configuration Summary appears orange for that line of the management group. Exceeding 16 storage systems in a cluster triggers a warning by turning that line red. As soon as the total number of storage systems reduces below the boundary, the summary bar returns to the previous indicator, either orange or green.

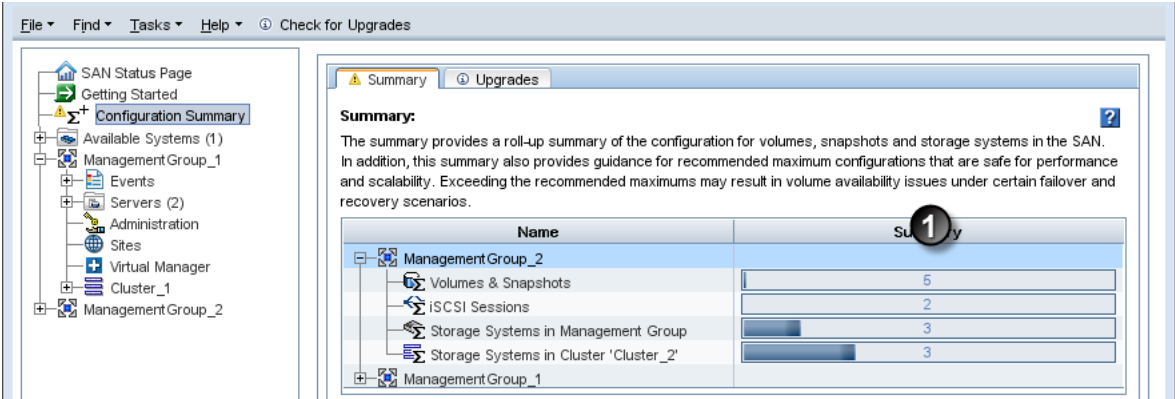
Reading the configuration summary

Each management group in the SAN is listed on the Configuration Summary. Underneath each management group is a list of the storage items tracked, such as storage systems, volumes, or iSCSI sessions. As items are added to the management group, the Summary graph fills in and the count is displayed in the graph. The Summary graph fills in proportionally to the optimum number for that item in a management group, as described in the “Best practices” (page 109).

Optimal configurations

Optimal configurations are indicated in green. For example, in Figure 45 (page 110), there are 15 storage systems in the management group “CJS1.” Those 15 storage systems are divided among the clusters “c” “c2” and “c3.” The length of the graph is relative to the recommended maximums in each category. For example, 3 storage systems in cluster c3 are closer to the cluster recommended maximum for storage systems than the 43 iSCSI sessions are to the maximum recommended iSCSI sessions for a management group.

Figure 45 Summary graph

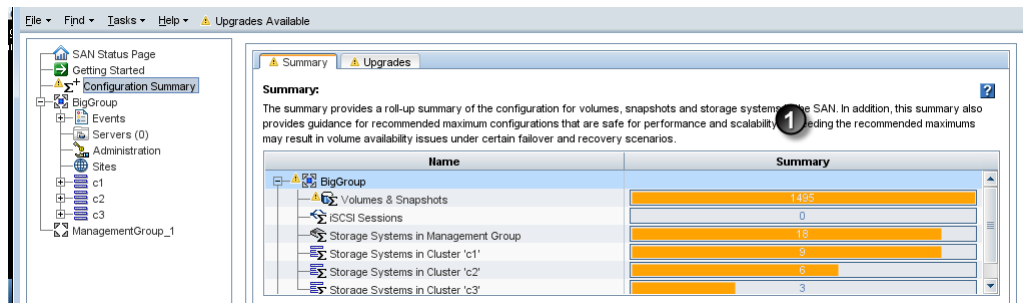


1. The items in the management group are all within optimum limits. The display is proportional to the optimum limits.

Configuration warnings

When any item nears a recommended maximum, it turns orange, and remains orange until the number is reduced to the optimal range. See Figure 46 (page 111).

Figure 46 Warning when items in the management group are reaching safe limits

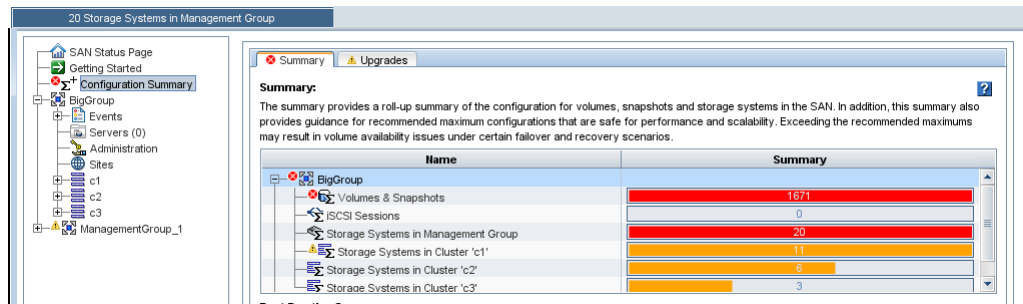


1. Volumes and snapshots are nearing the optimum limit. One cluster is nearing the optimum limit for storage systems.

Configuration errors

When any item exceeds a recommended maximum, it turns red, and remains red until the number is reduced. See Figure 47 (page 111).

Figure 47 Error when some item in the management group has reached its limit

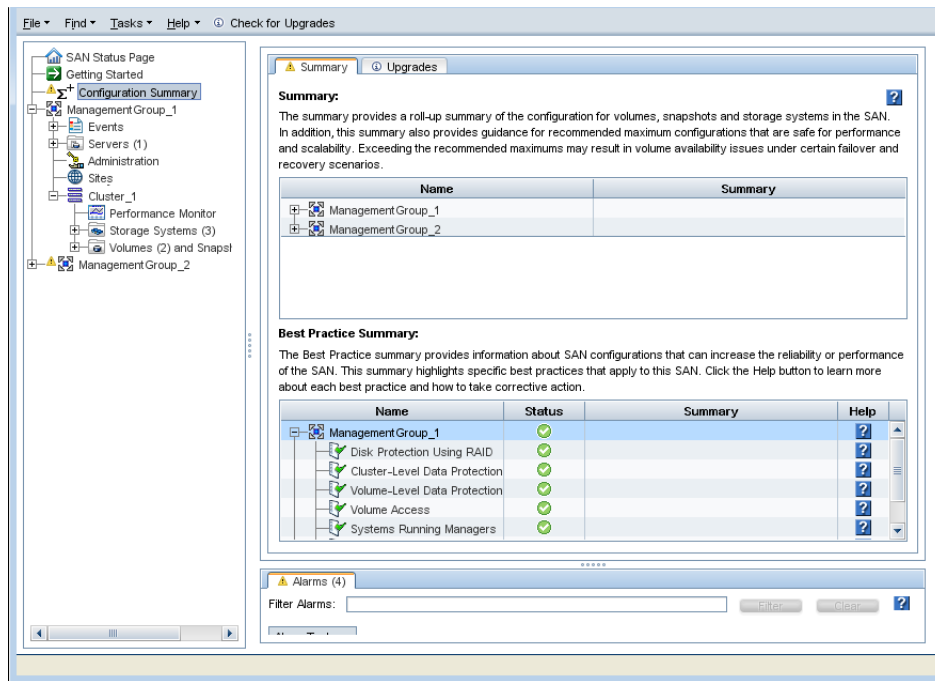


1. Volumes and snapshots have exceeded recommended maximums. One cluster remains near optimum limit.

Best Practice summary overview

The Best Practice summary provides an easy-to-use reference about best practices that can increase the reliability and/or performance of your SAN configurations. The Best Practice summary is available when you create a management group, and is found underneath the Configuration Summary. If you have more than one management group on the SAN, each group is listed in the summary.

Figure 48 Best Practice Summary for well-configured SAN



Expand the management group in the summary to see the individual categories that have recommended best practices. The summary displays the status of each category and identifies any conditions that fall outside the best practice. Click on a row to see details about that item's best practice.

Disk level data protection

Disk level data protection indicates whether the storage system has an appropriate disk RAID level set. For more information about configuring disk RAID levels, see [“Planning the RAID configuration”](#) (page 33).

Disk protection using RAID

Data protection is provided on an individual storage system by configuring RAID at any level other than 0. For a description of RAID levels, see [“RAID Levels”](#) (page 31).

Large single-system SATA cluster

If you are using a single large SATA storage system in a cluster, data protection is provided by configuring RAID 6 on that system. In addition to redundancy during normal operation, RAID 6 further protects the RAID array against data loss during degraded mode by tolerating one additional drive failure during this vulnerable stage.

Disk RAID Consistency

The best results for availability, reliability, and performance are achieved by using the same disk RAID level for all storage systems in a cluster. Mixed RAID levels in clusters are allowed to support a variety of administrative tasks. However, using the same RAID settings on all storage systems in a cluster ensures optimum availability and performance. For more information, see [“Managers overview”](#) (page 106).

Cluster-level data protection

Clusters of two or more systems provide the highest data availability. Clustered storage systems create the storage capacity for data volumes. Clusters are recommended to contain between 2 and 10 storage systems. See [“Storage systems in the cluster”](#) (page 110).

Volume-level data protection

Use a data protection level greater than Network RAID-0 to ensure optimum data availability if a storage system fails. For information about data protection, see [“Planning data protection”](#) (page 141).

Volume access

Use iSCSI load balancing to ensure better performance and better utilization of cluster resources. For more information about iSCSI load balancing, see [“iSCSI load balancing”](#) (page 230).

Systems running managers

Use the recommended number and type of managers to ensure optimum availability of your management group and volumes. Configurations with 3 or 5 managers are recommended for most single site installations. Three or five storage systems running managers is optimum. For fewer storage systems, use a Failover Manager as the third manager if possible. The Virtual Manager is also available for specific configurations. For a detailed discussion of managers and quorum, see [“Managers and quorum”](#) (page 106).

Network bonding

Bonding the available NIC cards in each storage system improves SAN performance and reliability. In most instances, Adaptive Load Balancing is the recommended bond. See [“Best practices”](#) (page 56).

Network bond consistency

To ensure consistent failover characteristics and traffic distribution, use the same network bond type in all the storage systems in a cluster. Inconsistent network bonds in storage systems in the same cluster may cause inconsistent traffic across NICs, and unexpected failover behavior depending on which system in the cluster fails. For more information, see [“Configuring network interface bonds”](#) (page 55).

Network flow control consistency

Set the network flow control settings the same on every storage system in the cluster. Inconsistent network flow control settings in storage systems in the same cluster may impact the connectivity between the systems. For more information, see [“Changing NIC flow control”](#) (page 52).

Network frame size consistency

Network frame size affects management data traffic and replication. Use the same network frame size in all storage systems in a cluster to ensure consistent data traffic and replication. Inconsistent network frame size may impact data traffic and replication in the cluster. For more information, see [“Changing NIC frame size”](#) (page 51).

Management group maintenance tasks

When you have an established management group, you may need to perform maintenance activities on the group.

Logging in to a management group

You must log in to a management group to administer the functions of that group.

1. In the navigation window, select a management group.
2. Log in by any of the following methods.
 - Double-click the management group.
 - Open the **Management Group Tasks** menu, and select **Log in to Management Group**. You can also open this menu from a right-click on the management group.
 - Click any of the **Log in to view** links on the Details tab.
3. Enter the user name and password, and click **Log In**.

When you log in to one storage system in a management group, you are logged in to all storage systems in that group.

Choosing which storage system to log in to

You can control which of the storage systems in a management group you log in to.

1. When the Log in to System window opens, click **Cancel**.

A message opens, asking if you want to log in to a different storage system.
2. Click **OK**.
3. The Log in to System window opens with a different storage system listed.
4. If that is the storage system you want, go ahead and log in. If you want to log in to a different storage system, repeat steps 1 and 2 until you see the storage system you want.

Logging out of a management group

Logging out of a management group prevents unauthorized access to that management group and the storage systems in that group.

1. In the navigation window, select a management group to log out of.
2. Click **Management Group Tasks** on the Details tab, and select **Log Out of Management Group**.

Adding a storage system to an existing management group

Storage systems can be added to management groups at any time. Add a storage system to a management group in preparation for adding it to a cluster.

1. In the navigation window, select an available storage system that you want to add to a management group.
2. Click **Storage System Tasks** on the Details tab, and select **Add to Existing Management Group**.
3. Select the desired management group from the drop-down list of existing management groups.
4. Click **Add**.
5. (Optional) If you want the storage system to run a manager, select the storage system in the management group, right-click, and select **Start Manager**.
6. Repeat steps 1 through 4 to add additional storage systems.

Starting and stopping managers

After adding the storage systems to the management group, start managers on the additional storage systems in the management group. The number of managers you start depends upon the overall design of your storage system. See [“Managers overview” \(page 106\)](#) for more information about how many managers to add.

Starting additional managers

1. In the navigation window select a storage system in the management group on which to start a manager.
2. Click **Storage System Tasks** on the Details tab, and select **Start Manager**.

Repeat these steps to start managers on additional storage systems.

Stopping managers

Under normal circumstances, you stop a manager when you are removing a storage system from a management group. You cannot stop the last manager in a management group. If you stop a manager that compromises fault tolerance, the Best Practice details identify an issue with systems that are running managers.

Deleting the management group is the only way to stop the last manager.

Implications of stopping managers

- Quorum of the storage systems may be decreased.
- Fewer copies of configuration data are maintained.
- Fault tolerance of the configuration data may be lost.
- Data integrity and availability may be compromised.



CAUTION: Stopping a manager can result in the loss of fault tolerance.

1. In the navigation window, select a management group, and log in.
2. Select the storage system for which you want to stop the manager.
3. Click **Storage System Tasks** on the Details tab, and select **Stop Manager**.
A confirmation message opens.
4. Click **OK** to confirm stopping the manager.

Editing a management group

Editing management group tasks include the following items:

- Changing local bandwidth priority.
- Editing Remote Bandwidth, done on the management group containing the remote snapshot. (See the *HP P4000 Remote Copy User Guide*, in “Chapter 2, Using Remote Copy,” the section about setting the remote bandwidth.)

Specialized editing tasks include:

- Disassociating management groups
- Setting Group Mode to normal

After making any changes to the management group, be sure to save the configuration data of the edited management group. See “[Saving management group configuration information](#)” (page 116).

Setting or changing the local bandwidth priority

After a management group has been created, edit the management group to change the local bandwidth priority. This is the maximum rate per second that a manager devotes to non-application processing, such as moving data. The default rate is 4 MB per second. You cannot set the range below 0.25 MB/sec.

Local bandwidth priority settings

The bandwidth setting is in MB per second. Use [Table 34 \(page 116\)](#) as a guide for setting the local bandwidth.

Table 34 Guide to local bandwidth priority settings

Network type	Throughput (MB/sec)	Throughput rating
Minimum	0.25	2 Mbps
Ethernet	1.25	10 Mbps
Factory Default	4.00	32 Mbps
Fast-Ethernet	12.50	100 Mbps
Half Gigabit-Ethernet	62.50	500 Mbps
Gigabit-Ethernet	128.00	1 Gbps
Bonded Gigabit-Ethernet (2)	256.00	2 Gbps
Bonded Gigabit-Ethernet (4)	512.00	4 Gbps

Set or change local bandwidth priority

1. In the navigation window, select a management group and log in.
2. Click **Management Group Tasks** on the Details tab, and select **Edit Management Group**.
3. Change the local bandwidth priority using the slider.
A default setting of 4, at the Application Access end of the slider, is more appropriate for everyday situations where many servers are busy with the volume. A setting of 40, at the Data Rebuild end of the slider, is most commonly used for quick data migration or copies when rebuilding or moving damaged volumes.
4. Click **OK**.
The new rate appears on the Details tab in the management group tab window.

Saving management group configuration information

Use View Management Group Configuration to save the configuration description. This creates a .txt file that lists the details of the management group configuration.

1. From the Tasks menu, select **Management Group**→**View Management Group Configuration**.
2. Select the management group if necessary.
3. Click Save in the Management Group Configuration window to save the configuration details in a .txt. file.

Safely shutting down a management group

Safely shut down a management group to ensure the safety of your data. Shutting down lets you:

- Perform maintenance on storage systems in that group
- Move storage systems around in your data center
- Perform maintenance on other equipment such as switches or UPS units
- Prepare for a pending natural disaster

Also, use a script to configure a safe shut down in the case of a controlled power down by a UPS. See [“Working with scripting” \(page 195\)](#). Sample scripts are available from the Customer Resource Center.

Shutting down a management group also relates to powering off individual storage systems and maintaining access to volumes. See the command line documentation, the *HP P4000 SAN/iQ Command Line Interface User Guide*, installed in the documentation directory of the program files.

Prerequisites

- Disconnect any hosts or servers that are accessing volumes in the management group.
- Wait for any restriping of volumes or snapshots to complete.
- Stop managers on storage systems that you plan to remove from the management group. You may want to add a FOM or start a Virtual Manager on a different storage system to maintain quorum and the best fault tolerance. See [“Stopping managers” \(page 115\)](#).

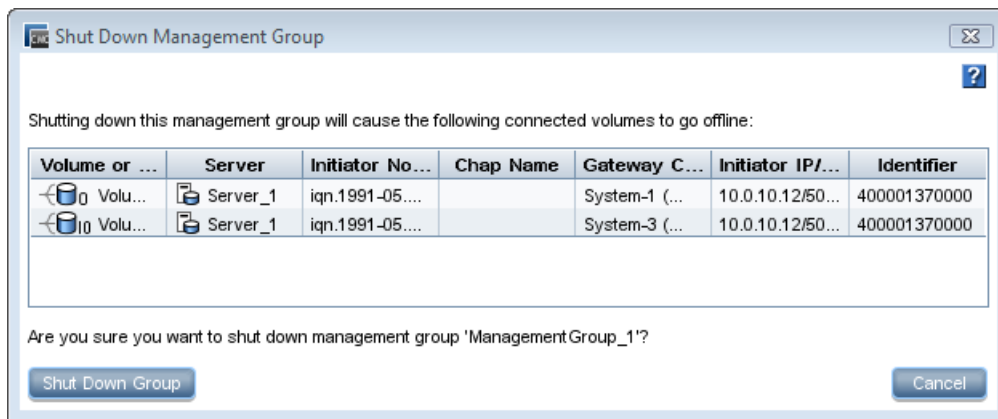
Shut down the management group

1. Log in to the management group that you want to shut down.
2. Click **Management Group Tasks** on the Details tab, and select **Shut Down Management Group**.
3. Click **Shut Down Group**.

If volumes are still connected to servers or hosts

After you click Shut Down Group, a confirmation window opens, listing volumes that are still connected and that will become unavailable if you continue shutting down the management group.

Figure 49 Notification of taking volumes offline



1. Stop server or host access to the volumes in the list.
2. Click **Shut Down Group**.

Start the management group back up

When you are ready to start up the management group, simply power on the storage systems for that group.

1. Power on the storage systems that were shut down.
2. Use Find in the CMC to discover the storage systems.

When the storage systems are all operating properly, the volumes become available and can be reconnected with the hosts or servers.

Restarted management group in maintenance mode

In certain cases the management group may start up in maintenance mode. Maintenance mode status usually indicates that either the management group is not completely restarted or the volumes are resynchronizing. When the management group is completely operational and the resynchronizing is complete, the management group status changes to normal mode.

Some situations which might cause a management group to start up in maintenance mode include the following:

- A storage system becomes unavailable, and the management group is shut down while that storage system is repaired or replaced. After the storage system is repaired or replaced and

the management group is started up, the management group remains in maintenance mode while the repaired or replaced storage system is resynchronizing with the rest of the management group.

- After a management group is shut down, a subset of storage systems is powered on. The management group remains in maintenance mode until the remaining storage systems are powered on and rediscovered in the CMC.
- For some reason, a storage system comes up, but it is not fully functional.

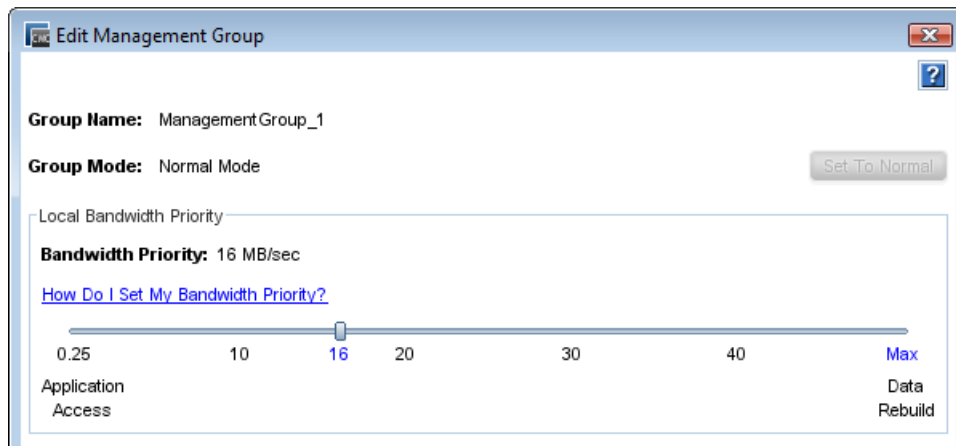
Manually change management group to normal mode

While the management group is in maintenance mode, volumes and snapshots are unavailable. You may get volumes and snapshots back online if you manually change the status from maintenance mode to normal mode, depending upon how your cluster and volumes are configured. However, manually changing from maintenance mode to normal mode causes the volumes in the management group to run in degraded mode while it continues resynchronizing, or until all the storage systems are up, or the reason for the problem is corrected.

CAUTION: If you are not certain that manually setting the management group to normal mode will bring your data online, or if it is not imperative to gain access to your data, do not change this setting.

1. In the navigation window, select the management group, and log in.
2. Click **Management Group Tasks** on the Details tab, and select **Edit Management Group**.

Figure 50 Manually setting management group to normal mode



3. Click **Set To Normal**.

Removing a storage system from a management group

When a storage system needs to be repaired or upgraded, remove it from the management group before beginning the repair or upgrade. Also remove a storage system from a management group if you are replacing it with another system.

Prerequisites

- Stop the manager on the storage system if it is running a manager. You may want to start a manager or Virtual Manager on a different storage system to maintain quorum and the best fault tolerance. See [“Stopping managers” \(page 115\)](#).
- (Optional) If the resultant number of storage systems in the cluster is fewer than the number required for the data protection level, you may have to change the data protection level of the volume(s) before removing the storage system from the cluster.

- Remove the storage system from the cluster. See [“Removing a storage system from a cluster” \(page 136\)](#).
- Let any restripe operations finish completely.

Remove the storage system

1. Log in to the management group from which you want to remove a storage system.
2. In the navigation window, select the storage system to remove.
3. Click **Storage System Tasks** on the Details tab, and select **Remove from Management Group**.
4. Click **OK** on the confirmation message.

In the navigation window, the storage system is removed from the management group and moved to Available Systems pool.

Deleting a management group

Delete a management group when you are completely reconfiguring your SAN and you intend to delete all data on the SAN.

⚠ CAUTION: When a management group is deleted, all data stored on storage systems in that management group is lost.

Prerequisites

- Log in to the management group.
- Remove all volumes and snapshots.
- Delete all clusters.

Delete the management group

1. In the navigation window, log in to the management group.
2. Click **Management Group Tasks** on the Details tab, and select **Delete Management Group**.
3. In the Delete Management Window, enter the management group name, and click **OK**.
After the management group is deleted, the storage systems return to the Available Systems pool.

Setting the management group version

If instructed by customer support, you can set the management group version back to a previous version of the software. Setting the version back requires using a special command line option before opening the CMC. Customer support will instruct you on using the command line option.

9 Using specialized managers

The SAN/iQ software provides two specialized managers that are used in specific situations. The Failover Manager is used in two-system and in Multi-Site SAN configurations to support automatic quorum management in the SAN. A virtual manager is added to a management group but is only started manually on a storage system if needed to maintain or regain quorum.

NOTE: You can only use one type of specialized manager in a management group. The Failover Manager is recommended for two-system configurations and in Multi-Site SAN configurations. The Virtual Manager is automatically added to two-system configurations if a Failover Manager is not added. However, the Failover Manager is the recommended solution for maintaining fault-tolerance on the SAN.

Definitions

Terms used in this chapter:

- **Failover Manager (FOM)**—A specialized manager running as a VMware guest operating system that can act as a quorum tie-breaker system when installed into a third location in the network to provide for automated failover/failback of the Multi-Site SAN clusters. It also added to two-system management groups to provide the optimum quorum configuration.
- **Virtual manager**—A manager which is added to a management group but is only started on a storage system if needed to regain or maintain quorum.
- **Regular manager**—A manager which is started on a storage system and operates according to the description of managers found in [“Managers overview” \(page 106\)](#).
- **Manager**—Any of these managers.

Failover Manager

The Failover Manager is a specialized version of the SAN/iQ software designed to run as a virtual appliance in either a VMware or Microsoft Hyper-V Server environment. The Failover Manager participates in the management group as a real manager in the system; however, it performs quorum operations only, not data movement operations. It is especially useful in a Multi-Site SAN configuration to manage quorum for the multi-site configuration without requiring additional storage systems to act as managers in the sites.

Planning the virtual network configuration

Before you install the Failover Manager on the network, plan the virtual network configuration, including the following areas:

- Design and configuration of the virtual switches and network adapters.
- Windows Server or ESX Server on which to install Failover Manager files, and a designated host name, and IP address.
- An iSCSI network for the virtual machine. The Failover Manager should be on the iSCSI network. If you do not have an existing virtual machine network configured on the iSCSI network/vswitch, create a new virtual machine network for the Failover Manager.

Failover Manager requirements

- Static IP address, or a reserved IP address if using DHCP.
- Server—not on the P4000 SAN—on which to install the Failover Manager.
- Only one Failover Manager per management group.
- Only a Failover Manager or a virtual manager in the a management group, not both.

- Failover Manager not inside a virtual Windows machine with VMware Server running.
- Failover Manager configured to start last after storage systems in the management group in server startup/shutdown list.

Using the Failover Manager on Microsoft Hyper-V Server

Install the Failover Manager on a Windows Server 2008 from the HP P4000 Management Software DVD, or from the DVD .iso image downloaded from the website:

<http://www.hp.com/go/P4000downloads>

The installer for the Failover Manager for Hyper-V Server includes a wizard that guides you through configuring the virtual machine on the network and powering on the Failover Manager.

Δ CAUTION: Do not install the Failover Manager on the HP P4000 SAN Solution, since this would defeat the purpose of the Failover Manager.

Minimum system requirements for using with Microsoft Hyper-V Server

- Supported versions of Windows Server
 - Windows Server 2008 R2 Standard
 - Windows Server 2008 R2 Enterprise
 - Windows Server 2008 R2 Datacenter
 - Windows Server 2008 R2 Server Core
 - Windows Server 2008 SP2, except for Core, which is not supported
- 15 GB available drive space on Windows server for the Failover Manager
- Hyper-V role installed and running.
- At least 1 GB of memory reserved for the Failover Manager.
- A single virtual CPU with at least 2000 MHz reserved.

Installing the Failover Manager for Hyper-V Server

1. Begin the installer one of the following ways:
 - Insert the HP P4000 Management Software DVD in the DVD drive.
 - Double-click the executable that you downloaded to start the installation.
2. Accept the terms of the License Agreement.
3. Choose a location for the Failover Manager virtual machine and a location for the virtual hard disks, and click **Next**.
4. Enter a name for the Failover Manager, select whether you want it powered on after it is installed, and click **Next**.
5. Enter the network information, including the host name and network addressing information, and click **Next**.
6. Finish the installation, reviewing the configuration summary, and click **Next**.

When the installer is finished, the Failover Manager is ready to be used in the HP P4000 SAN Solution

Next, use the **Find** function in the CMC to discover the Failover Manager, and then add it to a management group. See “Adding a storage system to an existing management group” (page 114) for more information.

Using the Failover Manager for VMware

Install the Failover Manager from the HP P4000 Management Software DVD, or from the DVD .iso image downloaded from the website:

<http://www.hp.com/go/P4000downloads>

The installer offers three choices for installing the Failover Manager for VMware.

- Failover Manager for ESX—The installer for the Failover Manager for ESX includes a wizard that guides you through configuring the virtual machine on the network and powering on the Failover Manager.
- Failover Manager for other VMware platforms—The installer for the Failover Manager for other VMware platforms includes a wizard that installs the Failover Manager. Then you find the Failover Manager, boot the system, and configure it for the P4000 SAN network.
- Failover Manager OVF files—If you choose to use the OVF package, you configure the data disk for the Failover Manager, and then configure the Failover Manager for the P4000 SAN network.

⚠ CAUTION: Do not install the Failover Manager on the HP P4000 SAN Solution, since this would defeat the purpose of the Failover Manager.

Minimum system requirements for using the Failover Manager with VMware ESX Server

- VMware ESX Server version 4.x or later.
- Microsoft .Net 2.0 or later on the installer client
- 1024 MB of RAM
- Bridged connection through the VMware Console, or assigned network in VI Client.
- 15 GB of disk space

Installing the Failover Manager for ESX Server

1. Begin the installer one of the following ways:
 - Insert the HP P4000 Management Software DVD in the DVD drive.
 - Double-click the executable that you downloaded to start the installation.
2. Accept the terms of the License Agreement.
3. Click **Install FOM for ESX** to begin the wizard.
4. In the Meta Launcher command line window that opens, enter **1** to run the installer CLI, or enter **2** to run the installer GUI. Both versions require the same information and perform the same installation.
 - If you selected the GUI installation in the Meta Launcher CLI window, the GUI installer window opens again and you click **Install FOM for ESX** again to begin the wizard.
5. Enter the login credentials for the ESX Server that will host the Failover Manager, and click **Next**.
6. Select the ESX Server host, verify the health status and configuration details, and click **Next**.
7. Select Failover Manager as the type of installation and click **Next**.
8. Select a datastore, not on the SAN, to store the virtual appliance files, and click **Next**.
9. Enter a name for the Failover Manager and enter the IP address or choose to use DHCP. Select the network interface to use and add multiple NICs if desired. Click **Next**.
10. Enter a name for the virtual machine, and whether to use VMDK or raw disk mapping, and click **Next**.

11. Select **No, I am done**, unless you plan to install another Failover Manager, and click **Next**.
If you want to install another Failover Manager, the wizard repeats the steps, using information you already entered, as appropriate.
12. Finish the installation, reviewing the configuration summary, and click **Deploy**.
When the installer is finished, the Failover Manager is ready to be used in the HP P4000 SAN Solution.

Next, use the Find function in the CMC to discover the Failover Manager, and then add it to a management group. See [“Adding a storage system to an existing management group”](#) (page 114) for more information.

Installing the Failover Manager using the OVF files with the VI Client

1. Begin the installer one of the following ways:
 - Insert the HP P4000 Management Software DVD in the DVD drive.
 - Double-click the executable that you downloaded to start the installation.
2. Click **Agree** to accept the terms of the License Agreement.
3. Click the link for **OVF files** to open a window from which you can copy the files to the ESX Server.

Configure the IP address and host name

1. In the inventory panel, select the new Failover Manager and power it on.
2. Select the Console tab and wait for the Failover Manager to boot.
3. When the Failover Manager finishes booting, enter **Start** and press **Enter** to log in to the Configuration Interface.
4. On the Configuration Interface main menu, tab to **Network TCP/IP Settings** and press **Enter**.
5. On the Available Network Devices window, tab to the network interface and press **Enter**.
6. On the Network Settings window, tab to the **Hostname** field and enter a host name for the Failover Manager. Use backspace to erase an entry if necessary.
This host name displays in the CMC only. It does not change the name of the original .VMX file or the name of the virtual machine in the VMware interface.
7. Tab to the method for setting the IP address.
If entering a static IP address, note that Gateway is a required field. If you do not have a gateway, enter 0.0.0.0.
8. Tab to **OK** and press **Enter**.
9. Press **Enter** again to confirm the action.
10. After the settings are configured, press **Enter** to confirm the IP address change.
11. On the Available Network Devices window, tab to **Back** and press **Enter**.
12. On the Configuration Interface, tab to **Log Out** and press **Enter**.

Finishing up with VI Client

1. In the VI Client Information Panel, click the **Summary** tab.
2. In the General section on the Summary tab, verify that the IP address and host name are correct, and that VMware Tools are running.

NOTE: If VMware Tools show “out of date” or “Unmanaged,” they are running correctly. The “out of date” or “Unmanaged” status is not a problem. VMware Tools are updated with each SAN/iQ software upgrade.

Next, use the Find function in the CMC to discover the Failover Manager, and then add it to a management group. See [“Adding a storage system to an existing management group”](#) (page 114) for more information.

Installing the Failover Manager for VMware Server or VMware Workstation

1. Begin the installer one of the following ways:
 - Insert the HP P4000 Management Software DVD in the DVD drive.
 - Double-click the executable that you downloaded to start the installation.
2. Click **Agree** to accept the terms of the License Agreement.
3. Click **Install FOM for other VMware platforms** to begin the wizard.
4. Accept the terms of the License Agreement, and click **Next**.
5. Accept the default location to install the files, or change the location, and click **Next**.
6. After the Failover Manager is installed, click **Finish** to complete the wizard.

Configuring the IP address and host name

1. In the VMware application, navigate to the location of the Failover Manager and start the virtual machine.
2. Wait for the Failover Manager to boot.
3. When the Failover Manager finishes booting, enter **Start** and press **Enter** to log in to the Configuration Interface.
4. On the Configuration Interface main menu, tab to **Network TCP/IP Settings** and press **Enter**.
5. On the Available Network Devices window, tab to the network interface and press **Enter**.
6. On the Network Settings window, tab to the **Hostname** field. Use backspace to erase an entry if necessary. Enter a host name for the Failover Manager.

This host name displays in the CMC only. It does not change the name of the original .VMX file or the name of the virtual machine in the VMware interface.
7. Tab to the method for setting the IP address.

If entering a static IP address, Gateway is a required field. If you do not have a gateway, enter 0.0.0.0.
8. Tab to **OK** and press **Enter**.
9. Press **Enter** again to confirm the action.
10. After the settings are configured, press **Enter** to confirm the IP address change.
11. On the Available Network Devices window, tab to **Back** and press **Enter**.
12. On the Configuration Interface, tab to **Log Out** and press **Enter**.

Next, use the Find function in the CMC to discover the Failover Manager, and then add it to a management group. See [“Adding a storage system to an existing management group”](#) (page 114) for more information.

Troubleshooting the Failover Manager on ESX Server

Use the following solutions to possible issues you encounter with the Failover Manager on an ESX Server.

Table 35 Troubleshooting for ESX Server installation

Issue	Solution
You want to reinstall the Failover Manager	<ol style="list-style-type: none">1. Close your CMC session.2. In the VI Client, power off the Failover Manager.3. Right-click, and select Delete from Disk.

Table 35 Troubleshooting for ESX Server installation *(continued)*

Issue	Solution
	<ol style="list-style-type: none"> 4. Copy fresh files into the virtual machine folder from the downloaded zip file or distribution media. 5. Open the VI Client, and begin again.
You cannot find the Failover Manager with the CMC, and cannot recall its IP address.	<ul style="list-style-type: none"> • The CMC displays the IP address of a system if it can be found. • Open a VI Client session and select the Summary tab for the system you want. The IP address and DNS name are displayed in the General information section.
In Linux	
If the installer does not start automatically	Run CMC_Installer.bin again.
In the VI Client	
You don't have the cursor available, or you don't have the keyboard available.	<ul style="list-style-type: none"> • If your cursor is missing, you are in console mode. Press Ctrl-Alt to regain the cursor. • If your keyboard is missing, move the mouse to the console window, and click once.
You want to see your Failover Manager, but the window is black.	Your console window has timed out. Click in the window with your mouse, and press any key.

Uninstalling the Failover Manager from VMware ESX Server

1. Remove the Failover Manager from the management group.
2. Power off the Failover Manager virtual machine in the VI Client.
3. Right-click the powered-off Failover Manager, and select **Delete from Disk**.

Virtual manager

A virtual manager is a manager that is added to a management group, but is not started on a storage system until it is needed to regain quorum. A virtual manager provides disaster recovery for one of two configurations:

- Configurations with only two storage systems. (A virtual manager will automatically be added when creating a management group using two storage systems.)
- Configurations in which a management group spans two geographic sites.

See [“Managers and quorum” \(page 106\)](#) for detailed information about quorum, fault tolerance, and the number of managers.

Because a virtual manager is available to maintain quorum in a management group when a storage system goes offline, it can also be used for maintaining quorum during maintenance procedures.

⚠ CAUTION: A virtual manager requires a manual intervention to recover quorum and can have undesirable effects when left running after quorum is recovered. That is why HP highly recommends that you use the Failover Manager rather than the virtual manager.

When to use a virtual manager

Use a virtual manager for disaster recovery in a two-site configuration, or a two-system configuration. You can also use a virtual manager to maintain quorum during storage system maintenance procedures, such as firmware upgrades.

Disaster recovery using a virtual manager

The virtual manager functions as an on-demand manager in a disaster-recovery situation. As an on-demand manager, it can be used to regain quorum and maintain access to data.

Management group across two sites with shared data

Using a virtual manager allows one site to continue operating if the other site fails. The virtual manager provides the ability to regain quorum in the operating site if one site becomes unavailable, or in one selected site if communication between the sites is lost. Such capability is necessary if volumes in the management group reside on storage systems in both locations.

Management group in a single location with two storage systems

If you create a management group with only two storage systems and do not have a Failover Manager, a virtual manager will be automatically added to the management group. The addition of the virtual manager provides the capability of regaining quorum if one manager becomes unavailable.

Storage system maintenance using a virtual manager

A virtual manager can also be used during maintenance to prevent loss of quorum. Adding a virtual manager to a management group enables you to start the virtual manager when you need to take a storage system offline for maintenance.

Requirements for using a virtual manager

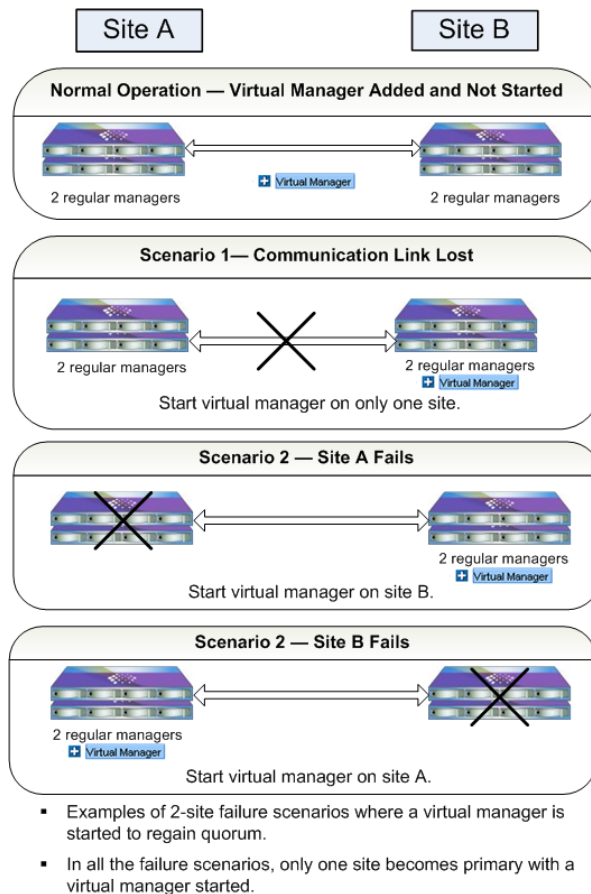
It is critical to use a virtual manager correctly. A virtual manager is added to the management group, but not started on a storage system until the management group experiences a failure and a loss of quorum. To regain quorum, you start the virtual manager on a storage system that is operating and in the site that is operational or primary.

Table 36 Requirements for using a virtual manager

Requirement	What it means		
Use a virtual manager with an even number of regular managers running on storage systems	Disaster recovery scenario	Number of regular managers running	Total number of managers, including the virtual manager
	Two separate sites with shared data	4	5
	Two storage systems in management group	2	3
Add a virtual manager when creating management group.	You cannot add a virtual manager after quorum has been lost. The virtual manager must be added to the management group before any failure occurs.		
A virtual manager must run only until the site is restored or communication is restored.	The virtual manager should run only until the site is restored and data is resynchronized, or until communication is restored and data is resynchronized.		

Correct uses of a virtual manager are shown in [Figure 51 \(page 127\)](#).

Figure 51 Two-site failure scenarios that are correctly using a virtual manager



Configuring a cluster for disaster recovery

In addition to using a virtual manager, you must configure your cluster and volumes correctly for disaster recovery. This section describes how to configure your system, including the virtual manager.

Best Practice

The following example describes configuring a management group with four storage systems in one cluster. The cluster spans two geographic sites with two storage systems at each site. The cluster contains a single volume with Network RAID-10 that spans both sites.

Configuration steps

The following configuration steps ensure that all the data is replicated at each site and the managers are configured correctly to handle disaster recovery.

1. Name storage systems with site-identifying host names.

To ensure that you can easily identify which storage systems reside at each site, use host names that identify the storage system location. See [“Changing the storage system hostname” \(page 22\)](#).

Management Group Name—**ManagementGroup_2**

Storage System Names

- Boulder-1
- Golden-1
- Boulder-2
- Golden-2

2. Create management group—plan the managers and virtual manager.

When you create the management group in the two-site scenario, plan to start two managers per site and add a virtual manager to the management group. You now have five managers for fault tolerance. See “Managers overview” (page 106).

3. Reorder the storage systems in the cluster alternating by site.

Create the cluster. Reorder the storage systems to the cluster in alternating order by site, as shown in the bulleted list. The order in which the storage systems are configured in the cluster determines the order in which copies of data are written to the volume. Alternating the storage systems by site location ensures that data is written to each site as part of the Network RAID-10 configured when you create the volume. See “Creating a cluster” (page 132).

Cluster Name—Cluster1

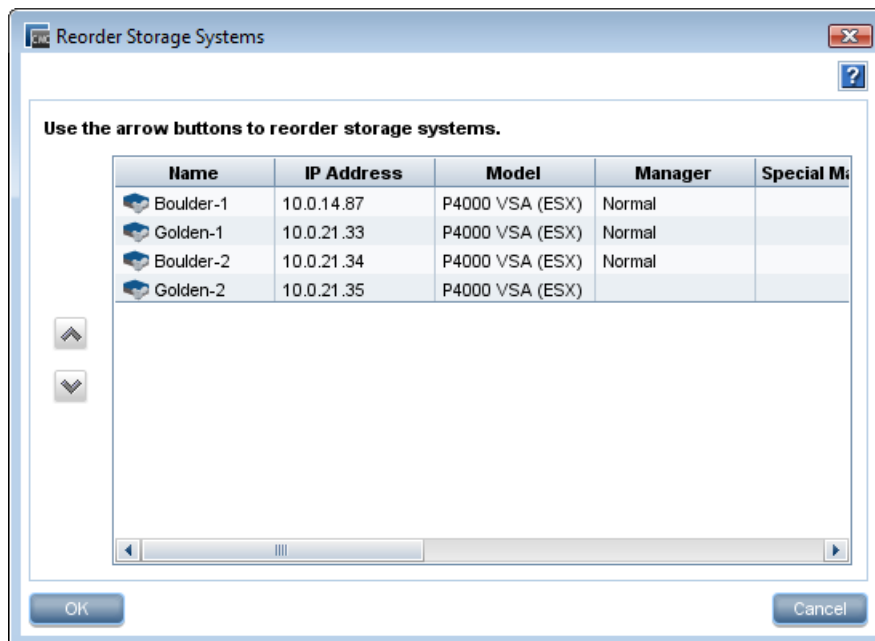
Ensure that the storage systems are in the cluster in the following order, as shown in Figure 52 (page 128).

- 1st storage system—Boulder-1
- 2nd storage system—Golden-1
- 3rd storage system—Boulder-2
- 4th storage system—Golden-2



CAUTION: If storage systems are listed in the cluster in any order other than alternating order by site, you will not have a complete copy of data on each site.

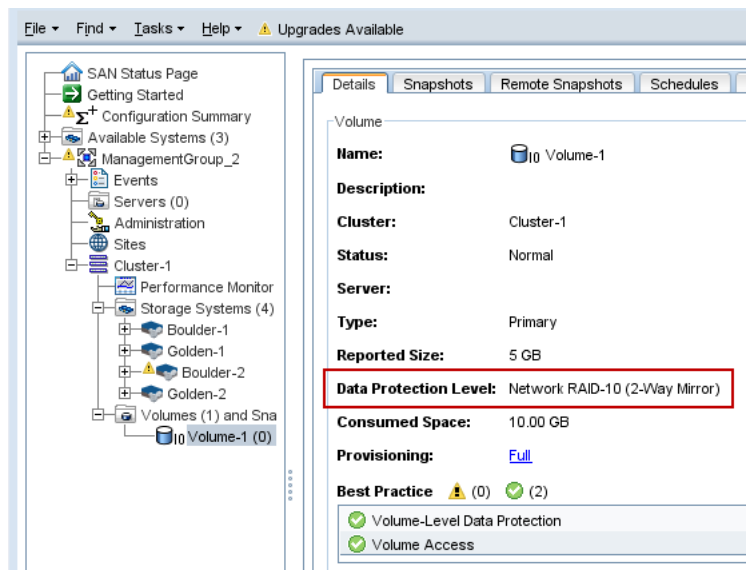
Figure 52 Moving storage systems in the cluster to alternating site order



4. Create the volume with Network RAID-10.

Network RAID-10 causes two copies of the data to be written to the volume. Because you added the storage systems to the cluster in alternating order, a complete copy of the data exists on each site. See “Planning data protection” (page 141).

Figure 53 Network RAID-10 volume on two-site cluster



Adding a virtual manager

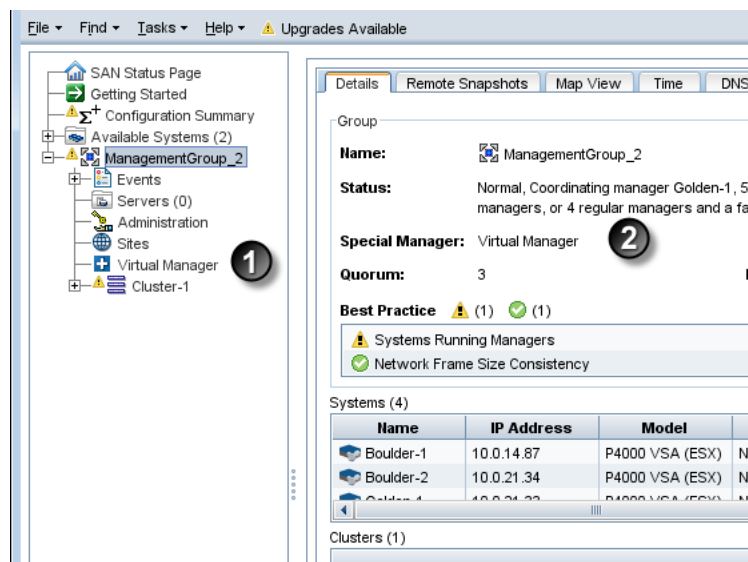
If a two-system management group is created without a Failover Manager, a virtual manager is added automatically to the group to ensure fault-tolerance.

NOTE: Best practice is to use a Failover Manager for a two-system management group.

1. Select the management group in the navigation window and log in.
2. Click **Management Group Tasks** on the Details tab, and select **Add virtual manager**.
3. Click **OK** to confirm the action.

The virtual manager is added to the management group. The Details tab lists the virtual manager as added, and the virtual manager icon appears in the management group.

Figure 54 Management group with virtual manager added



1. Virtual manager added

Starting a virtual manager to regain quorum

Only start a virtual manager when it is needed to regain quorum in a management group. “Two-site failure scenarios that are correctly using a virtual manager” (page 127) illustrates the correct way to start a virtual manager when necessary to regain quorum.

Two-site scenario, one site becomes unavailable

For example, in the two-site disaster recovery model, one of the sites becomes unavailable. On the site that remains up, all managers must be running. Select one of the storage systems at that site and start the virtual manager on it. That site then regains quorum and can continue to operate until the other site recovers. When the other site recovers, the managers in both sites re-establish communication and ensure that the data in both sites is resynchronized. When the data is resynchronized, stop the virtual manager to return to the disaster-tolerant configuration.

NOTE: If the unavailable site is not recoverable, you can create a new site with new storage systems and reconstruct the cluster. Contact Customer Support for help with cluster recovery. You must have the serial number of one of your storage systems when making a support call.

Two-site scenario, communication between the sites is lost

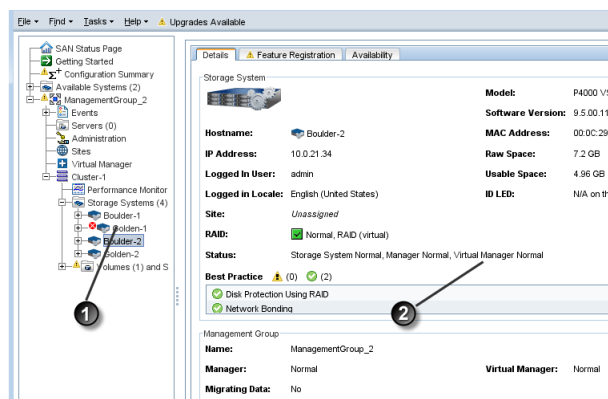
In this scenario, the sites are both operating independently. On the appropriate site, depending upon your configuration, select one of the storage systems, and start the virtual manager on it. That site then recovers quorum and operates as the primary site. When communication between the sites is restored, the managers in both sites reestablish communication and ensure that the data in both sites is resynchronized. When the data is resynchronized, stop the virtual manager to return to the disaster-tolerant configuration.

To start a virtual manager

A virtual manager must be started on a storage system, ideally one that isn’t already running a manager. However, if necessary, you can start a virtual manager on a storage system that is already running a manager, as shown in Figure 55 (page 130).

1. Select the storage system on which you want to start the virtual manager.
2. Click **Storage System Tasks** on the Details tab, and select **Start Virtual Manager**.

Figure 55 Starting a virtual manager when storage system running a manager becomes unavailable



1. Unavailable Manager
2. Virtual manager started

NOTE: If you attempt to start a virtual manager on a storage system that appears to be normal in the CMC, and you receive a message that the storage system is unavailable, start the virtual manager on a different storage system. This situation can occur when quorum is lost, because the CMC may still display the storage system in a normal state, even though the storage system is unavailable.

Verifying virtual manager status

Verify whether a virtual manager has been started, and if so, which storage system it is started on.

Select the virtual manager icon in the navigation window. The Details tab displays the location and status of the virtual manager.

Stopping a virtual manager

When the situation requiring the virtual manager is resolved—either the unavailable site recovers or the communication link is restored—you must stop the virtual manager. Stopping the virtual manager returns the management group to a fault tolerant configuration.

1. Select the storage system with the virtual manager.
2. Click **Storage System Tasks** on the Details tab, and select **Stop Virtual Manager**.
3. Click **OK** to confirm the action.

The virtual manager is stopped. However, it remains part of the management group and part of the quorum.

Removing a virtual manager from a management group

1. Log into the management group from which you want to remove the virtual manager.
2. Click **Management Group Tasks** on the Details tab, and select **Delete Virtual Manager**.
3. Click **OK** to confirm the action.

NOTE: The CMC will not allow you to delete a virtual manager if that deletion causes a loss of quorum.

10 Working with clusters

Clusters are groups of storage systems created in a management group. Clusters create a pool of storage from which to create volumes. The volumes seamlessly span the storage systems in the cluster. Expand the capacity of the storage pool by adding storage systems to the cluster.

Prerequisites

- An existing management group
- At least one storage system in the management group that is not already in a cluster
- A designated VIP address for the cluster

VIPs are required for iSCSI load balancing and fault tolerance, and for using HP DSM for MPIO. For more information, see [“iSCSI and the HP P4000 SAN Solution” \(page 229\)](#).

Clusters and storage systems

When creating clusters and when adding storage systems to clusters, consider storage system capacity and the number of storage systems planned.

- **Storage system capacity**—Typically you create clusters with storage systems of the same capacity. While clusters can contain storage systems with different capacities, all storage systems in a cluster operate at a capacity equal to that of the smallest-capacity storage system. The capacity limitation applies to the available capacity as determined by the disk RAID level, not the raw disk capacity. While you can mix storage systems with different RAID levels in a cluster, the capacity will be determined by the mix of RAID levels.

This capacity impact is particularly important to consider when planning to add additional storage systems to an existing cluster. If you add a storage system with a smaller capacity, the capacity of the entire cluster will be reduced. For example, If you have three storage systems, two of which have a RAID-determined capacity of 1 TB, and one of which has a RAID-determined capacity of 2 TB, all three storage systems operate at the 1 TB capacity.

- **Number of storage systems in clusters**—The optimum number of storage systems in a cluster ranges up to 10. For more information about the recommended maximum number of storage systems that can be added safely to a cluster, see [“Configuration Summary overview” \(page 108\)](#) or [“Working with management groups” \(page 103\)](#).

Creating a cluster

Creating a management group using the Management Group, Clusters, and Volumes wizard creates the first cluster in that management group. Use the following steps to create additional clusters in existing management groups.

1. Log in to the management group in which to create a cluster.
2. Right-click the storage system, and select **Add to New Cluster**.
3. Select **New Cluster**→**Standard Cluster**, and click **Next**.
4. Enter a meaningful name for the cluster.

A cluster name is case sensitive and must be from 1 to 127 characters.

5. Select one or more storage systems from the list to include in the cluster, and click **Next**.
6. Click **Add** in the Assign Virtual IPs and Subnet Masks window to add the VIP for the new cluster.
7. Enter the VIP and the Subnet Mask and click **OK**, and then click **Next**.
8. (Optional) Enter information for creating the first volume in the cluster, or select **Skip Volume Creation**.

9. Click **Finish**.
10. Review the summary information in the next window. Click **Close** when finished.

Adding an iSNS server (option)

Add an iSNS server.

NOTE: If using an iSNS server, you may not need to add Target Portals in the Microsoft iSCSI Initiator.

1. Select the new cluster in the tree. Click **iSCSI Tasks** menu, and select **Add iSNS Server**.
2. Right-click on the cluster or click **Cluster Tasks** and select Edit Cluster.
3. Select **Edit Cluster**→**Edit iSNS Servers**.
4. In the Edit iSNS Server window, click **Add Server**.
5. Enter the IP address of the iSNS server.
6. Click **OK**.
7. Click **OK** when finished.

Cluster Map View

After creating clusters and volumes and finish setting up the SAN, use the Map View tab for viewing the relationships between clusters, sites, volumes and systems. For more information on using the map view tools, see [“Using the display tools” \(page 14\)](#).

Monitoring cluster usage

The Use Summary, Volume Use, and System Use tabs provide detailed information about provisioning of volumes and snapshots and space usage in the cluster. See [“Ongoing capacity management” \(page 147\)](#) for information about the information reported on these tabs.

Editing a cluster

Editing a cluster comprises a group of tasks that include managing the storage systems in a cluster as well as changing a VIP, an iSNS server, or the cluster description.

Editing cluster properties

Change the cluster name or description.

1. Right-click the cluster or click **Cluster Tasks**.
2. Select **Edit Cluster**→**Edit Cluster Properties**.
3. In the Edit Cluster Properties change the name, or add or change the cluster description.
4. Click **OK** to save.

Editing iSNS servers

Add or change the iSNS servers for a cluster. If you change the IP address of an iSNS server, or remove the server, you may need to change the configuration that clients are using. Therefore, you may need to disconnect any clients before making this change.

Preparing clients

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the iSCSI initiator for those volumes.

Editing iSNS servers from the Cluster Tasks menu

1. Right-click the cluster or click **Cluster Tasks**.
2. Select **Edit Cluster**→**Edit iSNS Servers**.

3. In the Edit iSNS Servers window, select the VIP to change or delete, or click **Add** to add a new VIP.
4. Make the necessary changes and click **OK** to save.

Editing iSNS servers from the iSCSI Tasks menu

1. Select the cluster and click the **iSCSI** tab to bring it to the front.
2. Select the iSNS server IP to change and click **iSCSI Tasks**.
3. Select the appropriate menu item and make the desired changes.
4. Click **OK** to save.

Editing cluster VIP addresses

Anytime you add, change or remove the virtual IP address for iSCSI volumes, you are changing the configuration that servers are using. You should rebalance the iSCSI sessions after making the change. A command for rebalancing iSCSI sessions is available through the HP P4000 Command-Line Interface, CLIQ. The command syntax is in the *HP P4000 SAN/iQ Command-Line Interface User Guide*, which is installed with the CLI.

Preparing servers

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the iSCSI initiator for those volumes.

Editing cluster VIPs from the Cluster Tasks menu

Add, change, or delete the VIPs for a cluster, usually as part of a network reconfiguration.

1. Right-click the cluster or click **Cluster Tasks**.
2. Select **Edit Cluster**→**Edit Virtual IP addresses**.
3. In the Edit Virtual IPs window, select the VIP to change or delete, or click **Add** to add a new VIP.
4. Make the necessary changes and click **OK** to save.

Editing cluster VIP addresses from the iSCSI Tasks menu

1. Select the cluster and click the **iSCSI** tab to bring it to the front.
2. Select the VIP to change and click **iSCSI Tasks**.
3. Select the appropriate menu item for your task and make the desired changes.
4. Click **OK** to save.

Removing the virtual IP address

You can only remove a VIP if there is more than one VIP assigned to the cluster.

1. In the Edit Cluster window, click the **iSCSI** tab.
2. Select the VIP and click **Delete**.
3. Click **OK** to confirm the deletion.

Reconnecting volumes and applications after changing VIPs or iSNS servers

1. Reconfigure the iSCSI initiator with the changes.
2. Reconnect to the volumes.
3. Restart the applications that use the volumes.

Maintaining storage systems in clusters

Use the Edit Cluster menu to perform cluster maintenance tasks such as increasing capacity by adding a storage system, upgrading the cluster by swapping in new storage systems, or updating cluster VIPs and iSNS servers when necessary.

Adding a new storage system to a cluster

Add a new storage system to an existing cluster to expand the storage for that cluster. If the cluster contains a single storage system, adding a second storage system triggers a change for the volumes in the cluster to go from Network RAID 0 to Network RAID 10, which offers better data protection and volume availability. The window which opens describes the data protection offered by this change, and allows you to override the change for selected volumes, if necessary.

NOTE: Adding a storage system to a cluster causes a restripe of the data in that cluster. A restripe may take several hours or longer.

Adding a new storage system is not the same as replacing a repaired storage system with a new one. If you have repaired a storage system and want to replace it in the cluster, see [“Repairing a storage system” \(page 138\)](#).

Prerequisite

First add the storage system to the management group that contains the cluster.

To add a storage system

1. Select the cluster in the navigation window.
2. Click **Cluster Tasks**, and select **Edit Cluster**→**Add Storage Systems**.
If there are no storage systems in the management group available to add to the cluster, the Add Storage Systems menu item will be disabled.
3. Select one or more storage systems from the list.
4. Click **OK**.
5. Click **OK** again in the Edit Clusters window.
A confirmation message opens, describing the restripe that happens when a storage system is added to a cluster.
6. Click **OK** to confirm adding the storage system to the cluster.

Upgrading the storage systems in a cluster using cluster swap

Use cluster swap to upgrade all the systems in a cluster at one time. Using cluster swap allows upgrading the cluster with only one restripe of the data.

Prerequisite

Add the storage systems to swap to the management group that contains the cluster.

To swap storage systems

1. Select the cluster in the navigation window.
2. Click **Cluster Tasks**, and select **Edit Cluster**→**Swap Storage Systems**.
3. Select from the list all the storage systems to remove from the cluster and click **Remove Storage Systems**.
The selected storage systems disappear from the list.
4. Click **Add Storage Systems**.
The Add Storage Systems to Swap window opens with all the available storage systems in the management group.

5. Select the storage systems to swap into the cluster and click **OK**.

6. Click **OK** to complete the swap.

The swap operation may take some time, depending upon the number of storage systems swapped and the amount of data being restriped.

Reordering storage systems in a cluster

Reorder the systems in a cluster to control the stripe patterns, especially in a multi-site cluster.

1. Select the cluster in the navigation window.

2. Click **Cluster Tasks**, and select **Edit Cluster**→**Reorder Storage Systems**.

3. In the Reorder Storage Systems window, select a storage system and click the up or down arrow to move it to the desired position.

4. Click **OK** when the storage systems are in the desired order.

Exchange a storage system in a cluster

Use the Exchange storage system when you are ready to return a repaired storage system to the cluster. Exchanging a storage system is preferred to simply adding the repaired storage system to the cluster. Exchanging the storage system requires only a resync of the data in the cluster, rather than a restripe, which minimizes the time required to bring the cluster back to full operation. See [“How repair storage system works” \(page 138\)](#).

1. Select the cluster in the navigation window.

2. Click **Cluster Tasks**, and select **Edit Cluster**→**Exchange Storage Systems**.

3. In the Exchange Cluster Storage Systems window, select from the list the storage system to exchange and click **Exchange Storage Systems**.

4. Select the storage system from the list and click **OK**.

5. Click **OK** to complete the exchange.

Removing a storage system from a cluster

You can remove a storage system from a cluster only if the cluster contains sufficient storage systems to maintain the existing volumes and their data protection level. See [“Guide for volumes” \(page 154\)](#) for more information about editing volumes.

1. In the Edit Cluster window, select a storage system from the list.

2. Click **Remove Systems**.

In the navigation window, that storage system moves out of the cluster, but remains in the management group.

3. Click **OK** when you are finished.

NOTE: Removing a storage system causes a full cluster restripe.

Troubleshooting a cluster

Auto Performance Protection monitors individual storage system health related to performance issues that affect the volumes in the cluster.

Repairing a storage system provides a way to replace a failed disk in a storage system and minimize the time required to bring the storage system back to normal operation in the cluster with fully synchronized data.

Auto Performance Protection

If you notice performance issues in a cluster, a particular storage system may be experiencing slow I/O performance, overload, or latency issues. Identify whether Auto Performance Protection is operating by checking the storage system status on the storage system Details tab.

Auto Performance Protection is indicated by two unique statuses reported on the Details tab. You will also receive event notifications on these statuses. For information about setting up event notification, see [“Configuring events” \(page 88\)](#).

- **Storage System Overloaded.** The Overloaded status indicates that operations to the storage system are completing too slowly. During the overloaded state, volume availability is maintained while the storage system is quarantined in the cluster. While the storage system is quarantined it does not participate in I/O, which should relieve the performance degradation.

After the operations return to normal (in 10 minutes), the storage system is returned to active duty and resynced with the data that has changed since its quarantine. Volumes that depend on this storage system will then show “Resyncing” on the volume Details tab.

- **Storage System Inoperable.** The Inoperable status indicates that the storage system is unable to repair the slow I/Os, which may indicate a potential hardware problem. Volumes that depend on this storage system are unavailable. For information about how to determine volume availability, see the section [“Determining volume and snapshot availability” \(page 28\)](#).

Rebooting the storage system may return the status to Normal.

Auto Performance Protection and the VSA

The VSA will not report the Overloaded status, because there is no way to determine what may be affecting I/O on the underlying hardware. However, the VSA can accurately report when I/Os are not completing, and can return the Inoperable status.

Auto Performance Protection and other clusters


Auto Performance Protection operating on a storage system in one cluster will not affect performance for other clusters in the management group.

Checking storage system status

Easily identify whether Auto Performance Protection is active on a storage system in a cluster with performance issues.

1. Select the affected storage system in the navigation window.
The storage system icon will be blinking in the navigation tree.
2. Check the Status line on the Details tab.

Figure 56 Checking the storage system status on the Details tab



1. Status line

If status is Storage System Overloaded

Wait up to 10 minutes and check the status again. The status may return to Normal and the storage system will be resyncing.

If status is Storage System Inoperable

Reboot the storage system and see if it returns to Normal, when it comes back up.

If these statuses recur

This may be an indication that the underlying hardware problem still exists.

Repairing a storage system

Repairing a storage system allows you to replace a failed disk in a storage system that contains volumes configured for data protection levels other than Network RAID-0, and trigger only one resync of the data, rather than a complete restripe. Resyncing the data is a shorter operation than a restripe.

Prerequisites

- Volume must have Network RAID-10, Network RAID-10+1, Network RAID-10+2, Network RAID-5, or Network RAID-6.
- Storage system must have the blinking red and yellow triangle in the navigation window.
- If the storage system is running a manager, stopping that manager must not break quorum.

How repair storage system works

Using Repair Storage System to replace a failed disk includes the following steps:

- Using Repair Storage System from the Storage System Tasks menu to remove the storage system from the cluster
- Replacing the disk in the storage system
- Returning the storage system to the cluster


Because of the data protection level, removing and returning the storage system to the cluster would normally cause the remaining storage systems in the cluster to restripe the data twice—once when the storage system is removed from the cluster and once when it is returned.

The Repair Storage System command creates a placeholder in the cluster, in the form of a “ghost” storage system. This ghost storage system keeps the cluster intact while you remove the storage system, replace the disk, configure RAID, and return the storage system to the cluster. The returned storage system only has to resynchronize with the other two storage systems in the cluster.

Using the repair storage system command

When a storage system in a cluster has a disk failure, the navigation window displays the storage system and the cluster with a blinking triangle next to them in the tree. A disk inactive or disk off event appears in the Events list, and the Status label in the tab window shows the failure.

1. If the storage system is running a manager, stop the manager. See [“Stopping managers” \(page 115\)](#).
2. Right-click the storage system, and select **Repair Storage System**.

3. From the Repair Storage System window, select the item that describes the problem to solve. Click **More** for more detail about each selection.
 - Repair a disk problem
If the storage system has a bad disk, be sure to read [“Replacing a disk” \(page 42\)](#) before beginning the process.
 - Storage system problem
Select this choice if you have verified that the storage system must be removed from the management group to fix the problem. For more information about using Repair Storage System with a disk replacement, see [“Replacing disks” \(page 242\)](#).
 - Not sure
This choice offers the opportunity to confirm whether the storage system has a disk problem by opening the Disk Setup window so that you can verify disk status. As in the first choice, be sure to plan carefully for a disk replacement.
4. Click **OK**.
The storage system leaves the management group and moves to the Available Systems pool. A placeholder, or “ghost” storage system remains in the cluster. It is labeled with the IP address instead of the host name, and a special icon .
5. Replace the disk in the storage system and perform any other physical repairs.
Depending on the model, you may need to power on the disk and reconfigure RAID. See [“Replacing a disk” \(page 42\)](#).
6. Return the repaired storage system to the management group.
The ghost storage system remains in the cluster.

NOTE: The repaired storage system will be returned to the cluster in the same place it originally occupied to ensure that the cluster resyncs, rather than restripes. See [“Glossary” \(page 251\)](#) for definitions of restripe and resync.

7. [Optional] Start a the manager on the repaired storage system.

To return the repaired storage system to the cluster

Use the Exchange Storage System procedure to replace the ghost storage system with the repaired storage system. See [“Exchange a storage system in a cluster” \(page 136\)](#).

Deleting a cluster

Volumes and snapshots must be deleted or moved to a different cluster before deleting the cluster. For more information, see [“Deleting a volume” \(page 159\)](#) and [“Deleting a snapshot” \(page 175\)](#).

1. Log in to the management group that contains the cluster to delete.
2. In the navigation window, select the cluster to delete.
3. If there are any schedules to snapshot a volume or schedules to remote snapshot a volume for this cluster, delete them. See [“Deleting schedules to snapshot a volume” \(page 168\)](#).
4. From **Cluster Tasks**, select **Delete Cluster**.
A confirmation message opens. If the message says that the cluster is in use, you must delete the snapshots and volumes on the cluster.
 - a. Delete any volumes and snapshots, if necessary.

The cluster is deleted, and the storage systems return to the management group as available.

11 Provisioning storage

The SAN/iQ software uses volumes, including SmartClone volumes, and snapshots to provision storage to application servers and to back up data for recovery or other uses. Before you create volumes or configure schedules to snapshot a volume, plan the configuration you want for the volumes and snapshots.

Planning your storage configuration requires understanding how the capacity of the SAN is affected by the RAID level of the storage systems and the features of the SAN/iQ software. You also want to plan the level of data protection that is right for your configuration and storage requirements.

For example, if you are provisioning storage for MS Exchange, you will be planning the number and size of volumes you need for the databases and the log files. The capacity of the cluster that contains the volumes and snapshots is determined by the number of storage systems and the RAID level on them. The level of data protection takes into account the needs for data availability and data redundancy in your environment.

Understanding how the capacity of the SAN is used

The capacity of the SAN is a combination of factors.

- The first factor is the clustered capacity of the storage systems which is determined by the disk capacity and the RAID level.
See [“Planning the RAID configuration” \(page 33\)](#).
- The second factor is the effect of the data protection of the volumes and snapshots.
See [“Planning data protection” \(page 141\)](#).
- The third factor is the snapshot configuration, including schedules and retention policies.
See [“Managing capacity using volume size and snapshots” \(page 147\)](#).
- The fourth capacity factor is the impact of using Remote Copy as part of your backup and recovery strategy. Copying data to a remote cluster using remote snapshots, and then deleting that data from the application storage cluster, allows you to free up space on the application storage cluster more rapidly.

See the chapter [“Understanding and Planning Remote Copy”](#) in the *HP P4000 Remote Copy User Guide*.

Provisioning storage

Provisioning storage with the SAN/iQ software entails first deciding on the size of the volume presented to the operating system and to the applications. Next, decide on the configuration of snapshots, including schedules and retention policies.

Provisioning volumes

Configure volume size based on your data needs, how you plan to provision your volumes, and whether you plan to use snapshots. The SAN/iQ software offers both full and thin provisioning for volumes.

Table 37 Volume provisioning methods

Method	Settings
Full provisioning	Volume size x Network RAID level factor = amount of space allocated on the SAN
Thin provisioning	Volume size x Network RAID level factor >= amount of space allocated on the SAN

Full provisioning

Full provisioning reserves the same amount of space on the SAN as is presented to application servers. Full provisioning ensures that the application server will not fail a write. When a fully provisioned volume approaches capacity, you receive a warning that the disk is nearly full.

Thin provisioning

Thin provisioning reserves less space on the SAN than is presented to application servers. The SAN/iQ software allocates space as needed when data is written to the volume. Thin provisioning also allows storage clusters to provision more storage to application servers than physically exists in the cluster. When a cluster is over-provisioned, thin provisioning carries the risk that an application server will fail a write if the storage cluster has run out of disk space. The SAN/iQ software adds utilization and over-provisioned events as the cluster approaches 100% utilization. You can add capacity to the cluster or delete unneeded snapshots to accommodate additional volume growth.

NOTE: Paying attention to the space utilization events on over-provisioned storage clusters is critical to prevent a write failure on a thin volume.

Best practice for setting volume size

Create the volume with the size that you currently need. Later, if you need to make the volume bigger, increase the volume size in the CMC and then expand the disk on the server. In Microsoft Windows, you expand a basic disk using Windows Logical Disk Manager and Diskpart. For detailed instructions, see [“Changing the volume size on the server”](#) (page 152).

Planning data protection

Data protection results from creating data redundancy for volumes on the SAN. Configure data protection levels, called Network RAID, when you create a volume. Because data is stored redundantly on different storage systems, all data protection levels are tied to the number of available storage systems in a cluster. Network RAID achieves data protection using can use either replication or parity.

- Data protection with replication - Data protection using replication is achieved using Network RAID-10, Network RAID-10+1, or Network RAID-10+2, which store two, three, or four mirrored copies of the data.
- Data protection with parity - Data protection using parity is achieved using Network RAID-5 and Network RAID-6, which store data and parity dynamically according to the number of storage systems in a cluster. Under some workloads that write the data infrequently, Network RAID-5 and Network RAID-6 provide better capacity utilization and similar high availability to Network RAID-10 and Network RAID-10+1.

Former terminology in release 8.1 and earlier

Before release 8.5, you configured volumes with replication levels.

Volume Replication Level	Data Protection Level
None	Network RAID-0 (None)
2-Way replication	Network RAID-10 (2-Way Mirror)
3-Way replication	Network RAID-10+1 (3-Way Mirror)
4-Way replication	Network RAID-10+2 (4-Way Mirror)
–	Network RAID-5 (Single Parity) (new)
–	Network RAID-6 (Dual Parity) (new)

Data protection level

Seven data protection levels are available, depending upon the number of available storage systems in the cluster.

Table 38 Setting a data protection level for a volume

With this number of available storage systems in cluster	Select any of these data protection levels	For this number of copies
1	<ul style="list-style-type: none"> • Network RAID-0 (None) 	<ul style="list-style-type: none"> • One copy of data in the cluster.
2	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 (2-Way Mirror) 	<ul style="list-style-type: none"> • One copy of data in the cluster. • Two copies of data in the cluster.
3	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 (2-Way Mirror) • Network RAID-10+1 (3-Way Mirror) • Network RAID-5 (Single Parity) 	<ul style="list-style-type: none"> • One copy • Two copies • Three copies • Data striped on three storage systems including single parity.
At least 4	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 (2-Way Mirror) • Network RAID-10+1 (3-Way Mirror) • Network RAID-10+2 (4-Way Mirror) • Network RAID-5 (Single Parity) 	<ul style="list-style-type: none"> • One copy • Two copies • Three copies • Four copies • Data striped on four storage systems including single parity.
At least 5	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 • Network RAID-10+1 (3-Way Mirror) • Network RAID-10+2 (4-Way Mirror) • Network RAID-5 (Single Parity) • Network RAID-6 (Dual Parity) 	<ul style="list-style-type: none"> • One copy • Two copies • Three copies • Four copies • Data striped on five storage systems including single parity. • Data striped on five storage systems including dual parity.
At least 6	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 • Network RAID-10+1 (3-Way Mirror) • Network RAID-10+2 (4-Way Mirror) • Network RAID-5 (Single Parity) • Network RAID-6 (Dual Parity) 	<ul style="list-style-type: none"> • One copy • Two copies • Three copies • Four copies • Data striped on five storage systems including single parity. • Data striped on six storage systems including dual parity.
At least 7	<ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-10 • Network RAID-10+1 (3-Way Mirror) • Network RAID-10+2 (4-Way Mirror) • Network RAID-5 (Single Parity) • Network RAID-6 (Dual Parity) 	<ul style="list-style-type: none"> • One copy • Two copies • Three copies • Four copies • Data striped on five storage systems including single parity. • Data striped on seven storage systems including dual parity.

How data protection levels work

The system calculates the actual amount of storage resources needed for all data protection levels. When you choose Network RAID-10, Network RAID-10+1, or Network RAID-10+2, data is striped and mirrored across either two, three, or four adjacent storage systems in the cluster.

When you choose Network RAID-5 or Network RAID-6, the layout of the data stripe, including parity, depends on both the Network RAID mode and cluster size. These Network RAID stripes are rotated across the storage cluster to support configurations where the number of storage systems in the cluster is larger than the Network RAID stripe. When using Network RAID 5, the Network RAID stripe spans three to five storage systems including single parity. When using Network RAID 6, the Network RAID stripe spans five to seven storage systems including dual parity.

CAUTION: A management group with two storage systems and a Failover Manager is the minimum configuration for automated fault tolerant operation. Although the SAN/iQ software allows you to configure Network RAID-10 on two storage systems, this does not guarantee data availability in the event that one storage system becomes unavailable, due to the communication requirements between managers. See [“Managers overview” \(page 106\)](#).

CAUTION: Any volume with Network RAID-0 is not protected from complete system failure or reboot.


Network RAID-10 (2-Way Mirror)

Network RAID-10 data is striped and mirrored across two storage systems. Network RAID-10 is the default data protection level assigned when creating a volume, as long as there are two or more storage systems in the cluster. Data in a volume configured with Network RAID-10 is available and preserved in the event that one storage system becomes unavailable.

Network RAID-10 is generally the best choice for applications that write to the volume frequently and don't need to tolerate multiple storage system failures. Such applications include databases, email, and server virtualization. Network RAID-10 is also good for Multi-Site SANs. Using Network RAID-10 in a Multi-Site SAN ensures that data remains available in the event that one site becomes unavailable. However, if one site does go down, the Network RAID-10 volumes are then not protected from complete system failure or reboot.

[Figure 57 \(page 143\)](#) illustrates the write patterns on a cluster with four storage systems configured for Network RAID-10.

Figure 57 Write patterns in Network RAID-10 (2-Way Mirror)




Network RAID-10+1 (3-Way Mirror)

Network RAID-10+1 data is striped and mirrored across three or more storage systems. Data in a volume configured with Network RAID-10+1 is available and preserved in the event that any two storage systems become unavailable.

Best applications for Network RAID-10+1 are those that require data availability even if two storage systems in a cluster become unavailable.

Figure 58 (page 144) illustrates the write patterns on a cluster with four storage systems configured for Network RAID-10+1.

Figure 58 Write patterns in Network RAID-10+1 (3-Way Mirror)




Network RAID-10+2 (4-Way Mirror)

Network RAID-10+2 data is striped and mirrored across four or more storage systems. Data in a volume configured with Network RAID-10+2 is preserved in the event that any three storage systems become unavailable. Network RAID-10+2 is designed for Multi-Site SANs to preserve data in the event of an entire site becoming unavailable.

Best use for Network RAID-10+2 volumes is for data that must be synchronously replicated between two locations and that must remain fully redundant in the case of an entire site failure. Using Network RAID-10+2 ensures that data remains available after half of the SAN is unavailable, and continues to remain available even with the loss of a single storage system in the remaining site.

Figure 59 (page 144) illustrates the write patterns on a cluster with four storage systems configured for Network RAID-10+2.

Figure 59 Write patterns in Network RAID-10+2 (4-Way Mirror)



Network RAID-5 (Single Parity)


Network RAID-5 divides the data into stripes and adds parity. Network RAID stripe spans three to five storage systems including single parity. Data in a volume configured with Network RAID-5 is available and preserved in the event that any single storage system becomes unavailable.

Network RAID-5 volumes are configured as thin provisioned by default.

Best applications for using Network RAID-5 volumes include applications with mostly read, sequential workloads, such as file shares and archiving.

Figure 60 (page 145) illustrates the write patterns on a cluster with four storage systems configured for Network RAID-5.

Figure 60 Write patterns and parity in Network RAID-5 (Single Parity)



1. Parity for data blocks A, B, C
2. Parity for data blocks D, E, F
3. Parity for data blocks G, H, I
4. Parity for data blocks J, K, L

Network RAID-6 (Dual Parity)


Network RAID-6 divides the data into stripes and adds parity. When using Network RAID 6, the Network RAID stripe spans five to seven nodes including dual parity. Data in a volume configured with Network RAID-6 is available and preserved in the event that any two storage systems become unavailable.

Network RAID-6 volumes are configured as thin provisioned by default.

Best applications for using Network RAID-6 volumes include applications with mostly read, sequential workloads on larger clusters, such as file shares and archiving.

Figure 61 (page 146) illustrates the write patterns on a cluster with six storage systems configured for Network RAID-6.

Figure 61 Write patterns and parity in Network RAID-6 (Dual Parity)



1. P1 is parity for data blocks A, B, C, D
2. P2 is parity for data blocks E, F, G, H
3. P3 is parity for data blocks I, J, K, L
4. P4 is parity for data blocks M, N, O, P

Provisioning snapshots

Snapshots provide a copy of a volume for use with backup and other applications. You create snapshots from a volume on the cluster.

Snapshots are always thin provisioned. Thin provisioning snapshots saves actual space in the SAN, while letting you have more snapshots without the concern of running out of cluster space.

Snapshots can be used for multiple purposes, including:

- Source volumes for data mining and other data use
- Source volumes for creating backups
- Data or file system preservation before upgrading software
- Protection against data deletion
- File-level restore without tape or backup software

Snapshots versus backups

Backups are typically stored on different physical devices such as tapes. Snapshots are stored in the same cluster as the volume. Therefore, snapshots protect against data deletion, but not device or storage media failure. Use snapshots along with backups to improve your overall data backup strategy.

At any time you can roll back to a specific snapshot. When you do roll back, you must delete all the snapshots created after that snapshot. Also, using an iSCSI initiator, you can mount a snapshot to a different server and recover data from the snapshot to that server.

The effect of snapshots on cluster space

Snapshots take up space on the cluster. Because snapshots are a thin provisioned space, they save space compared to a full provisioned space.

Plan how you intend to use snapshots, and the schedule and retention policy for schedules to snapshot a volume. Snapshots record changes in data on the volume, so calculating the rate of changed data in the client applications is important for planning schedules to snapshot a volume.

NOTE: Volume size, provisioning, and using snapshots should be planned together. If you intend to use snapshots, review [“Using snapshots” \(page 161\)](#).

Managing capacity using volume size and snapshots

When you create a snapshot of a volume, the original volume is actually saved as the snapshot, and a new volume (the “writable” volume) with the original name is created to record any changes made to the volume’s data after the snapshot was created. Subsequent snapshots record only changes made to the volume since the previous snapshot. Snapshots are always created as a thin provisioned space, no matter whether its original volume is full or thin provisioned.

Volume size and snapshots

One implication of the relationship between volumes and snapshots is that the space used by the writable volume can become very small when it records only the changes that have occurred since the last snapshot was taken. This means that less space may be required for the writable volume.

Over time, you may find that space allocated for snapshots becomes larger, and the volume itself becomes relatively small.

Schedules to snapshot a volume and capacity

When you have schedules to snapshot a volume, the recurrence or frequency, and the retention policy for the schedules affect the amount of space used in the cluster. For example, it is possible for a new snapshot and one snapshot scheduled for deletion to coexist in the cluster for some period of time. If there is not sufficient room in the cluster for both snapshots, the scheduled snapshot will not be created, and the schedule will not continue until an existing snapshot is deleted. Therefore, if you want to retain (n) snapshots, the cluster should have space for (n+1).

Deleting snapshots

Another factor to note in planning capacity is the fact that when a snapshot is deleted, that snapshot’s data is added to the snapshot or volume directly above it (the next newer snapshot). The amount of space allocated for the volume or snapshot directly above the deleted snapshot increases. See [“Ongoing capacity management” \(page 147\)](#) for detailed information about reviewing capacity.

Ongoing capacity management

One of the critical functions of managing a SAN is monitoring usage and capacity. The CMC provides detailed information about overall cluster capacity and usage, as well as detail about provisioning and storage system capacity.

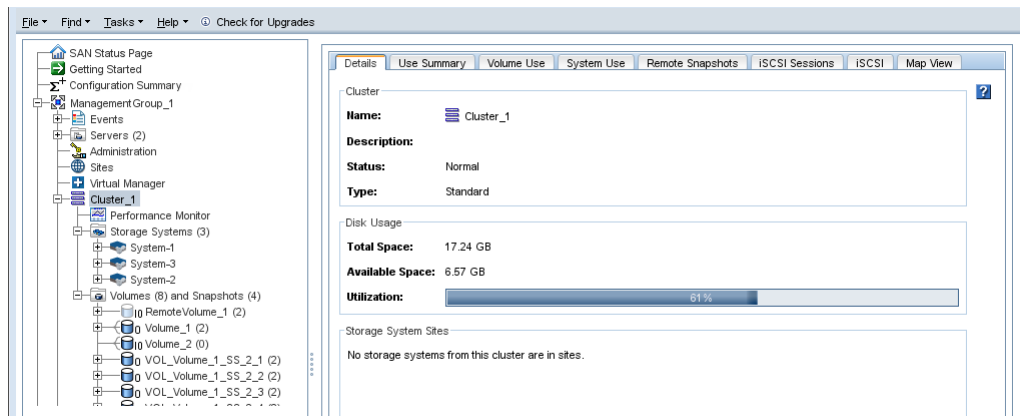
Number of volumes and snapshots

For information about the recommended maximum number of volumes and snapshots that can be created in a management group, see [“Configuration Summary overview” \(page 108\)](#).

Reviewing SAN capacity and usage

You can review detailed information about the capacity of your cluster, the volumes it contains, and the provisioning of the storage systems in the cluster. This information is presented in a series of tab windows displayed at the cluster level.

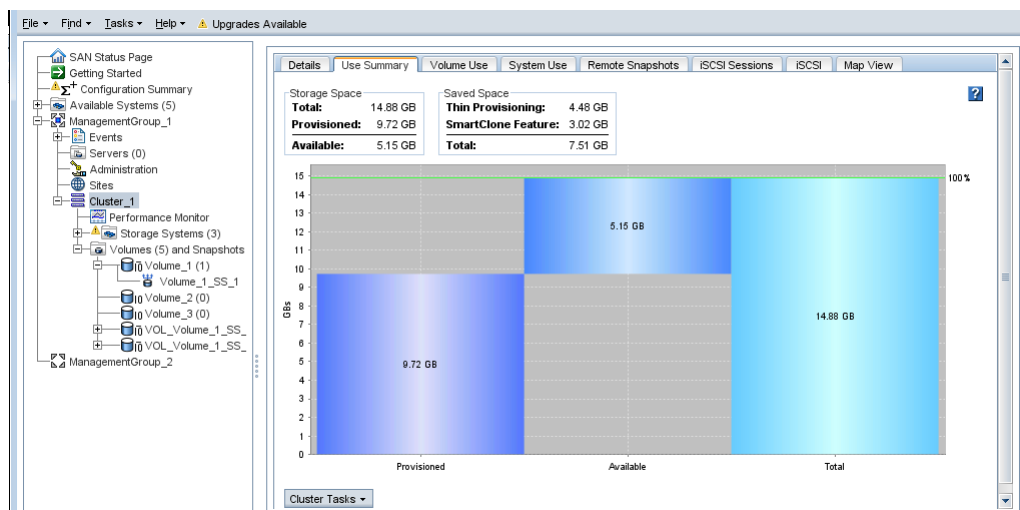
Figure 62 Cluster tab view



Cluster use summary

The Use Summary window presents information about the storage space available in the cluster.

Figure 63 Reviewing the Use Summary tab



In the Use Summary window, the Storage Space section lists the space available on the storage systems in the cluster. Saved space lists the space saved in the cluster by using thin provisioning and the SmartClone feature. The graphical display illustrates the amount of space in the cluster and its allocation as provisioned and available. Storage space is broken down as shown in [Table 39 \(page 148\)](#).

Table 39 Information on the Use Summary tab

Category	Description
Table information	
Storage Space	
Total	Combined space available in the cluster for storage volumes and snapshots.
Provisioned	Amount of space allocated for storage, including both volumes and snapshots. This value increases as snapshots are taken, or as a thinly provisioned volume grows.
Available	Amount of space remaining in the cluster that has not been allocated for storage. This value decreases as volumes and

Table 39 Information on the Use Summary tab *(continued)*

Category	Description
	snapshots are created, or as thinly provisioned volumes grow.
Saved Space	
Thin Provisioning	The space saved by thin provisioning volumes. This space is calculated by the system.
SmartClone Feature	Space saved by using SmartClone volumes is calculated using the amount of data in the clone point and any snapshots below the clone point. Only as data is added to an individual SmartClone volume does it consume space on the SAN.
Total	Approximate total amount of space saved by using thin provisioning and SmartClone volumes.
Graph information	
Provisioned	Amount of space allocated for volumes and snapshots.
Available	Amount of space remaining in the cluster that has not been allocated for storage. This value decreases as volumes and snapshots are created, or as thinly provisioned volumes grow.
Total	Combined space available in the cluster for storage volumes and snapshots.
Provisionable	This graph appears only if the cluster is over-provisioned. It displays the total space that volumes and snapshots can grow to fill, and this value can exceed the physical capacity of the SAN. As an over-provisioned cluster approaches the physical capacity of the SAN, a series of warnings displays on the Use Summary window.

Volume use summary

The Volume Use window presents detailed information about the volume characteristics that affect the utilization of the cluster.

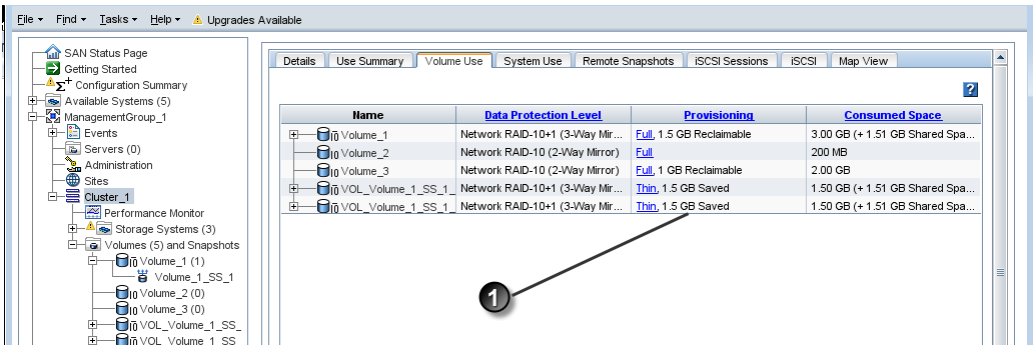
Table 40 Information on the Volume Use tab

Category	Description
Name	Name of the volume or snapshot.
Data Protection Level	The level of data protection configured for the volume. Snapshots inherit the data protection level of the parent volume.
Provisioning	<p>Volumes can be either full or thin provisioned. The Provisioning column also details space saving options for the different types of volumes you can create on the SAN, as shown in Figure 64 (page 150). The space calculations take into account both the type of volume and the data protection level. Use this information to help you manage space use on the SAN.</p> <ul style="list-style-type: none"> Thin provisioning saves space on the SAN by only allocating a fraction of the configured volume size. Therefore, the space saved on the SAN is reflected in this column. As data is added to the volume, thin provisioning grows the allocated space. You can expect

Table 40 Information on the Volume Use tab (continued)

Category	Description
	<p>to see the space saved number decrease as data on the volume increases.</p> <ul style="list-style-type: none">Full provisioning allocates the full amount of space for the size of the volume. Reclaimable space is the amount of space that you can get back if this fully provisioned volume is changed to thinly provisioned.
Consumed Space	<p>Amount of space used by actual data volumes or snapshots. The consumed space is the amount you can regain if you delete the volume and its snapshots. If the volume has SmartClone clone points below it, the clone point and any snapshot below it comprise the Shared Space listed next to the consumed space. To regain the Shared Space, you must delete the clone point and any snapshots below it. For instructions on deleting clone points, see "Deleting the clone point " (page 193)</p> <p>Deleting files or data from client applications does not decrease the used space. For more information, see "Measuring disk capacity and volume size" (page 151).</p>

Figure 64 Viewing the space saved or reclaimable in the Volume Use tab



1. Space saved or reclaimable displayed here

System Use summary

The System Use window presents a representation of the space provisioned on the storage systems in the cluster.

Figure 65 Viewing the System Use tab

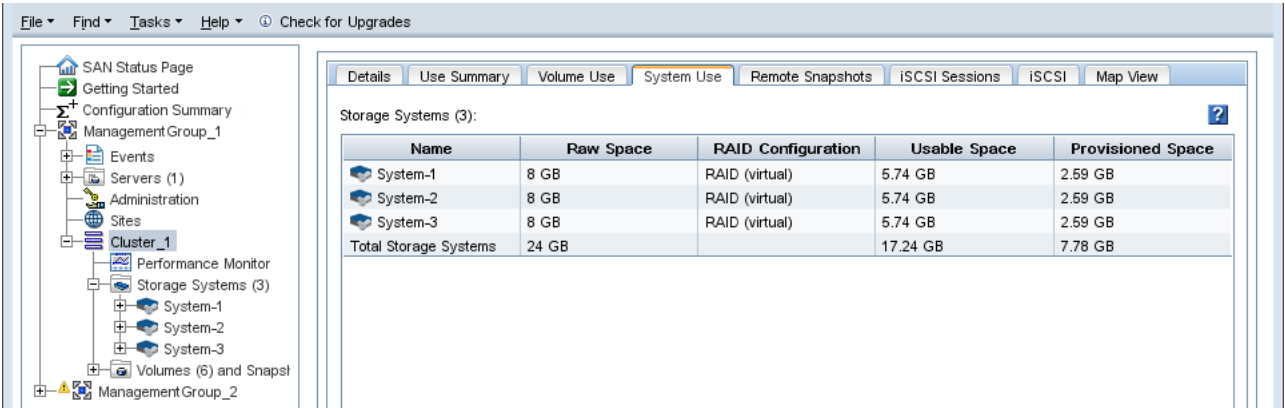


Table 41 Information on the System Use tab

Category	Description
Name	Host name of the storage system.
Raw space	Total amount of disk capacity on the storage system. Note: Storage systems with greater capacity will only operate to the capacity of the lowest capacity storage system in the cluster.
RAID configuration	RAID level configured on the storage system.
Usable space	Space available for storage after RAID has been configured.
Provisioned space	Amount of space allocated for volumes and snapshots.

Measuring disk capacity and volume size

All operating systems that are capable of connecting to the SAN via iSCSI interact with two disk space accounting systems—the block system and the native file system (on Windows, this is usually NTFS).

Table 42 Common native file systems

OS	File System Names
Windows	NTFS, FAT
Linux	EXT2, EXT3
Netware	NWFS
Solaris	UFS
VMWare	VMFS

Block systems and file systems

Operating systems see hard drives (both directly connected [DAS] and iSCSI connected [SAN]) as abstractions known as “block devices”: arbitrary arrays of storage space that can be read from and written to as needed.

Files on disks are handled by a different abstraction: the file system. File systems are placed on block devices. File systems are given authority over reads and writes to block devices.

iSCSI does not operate at the file system level of abstraction. Instead, it presents the iSCSI SAN volume to an OS such as Microsoft Windows as a block device. Typically, then, a file system is created on top of this block device so that it can be used for storage. In contrast, an Oracle database can use an iSCSI SAN volume as a raw block device.

Storing file system data on a block system

The Windows file system treats the iSCSI block device as simply another hard drive. That is, the block device is treated as an array of blocks which the file system can use for storing data. As the iSCSI initiator passes writes from the file system, the SAN/iQ software simply writes those blocks into the volume on the SAN. When you look at the CMC, the used space displayed is based on how many physical blocks have been written for this volume.

When you delete a file, typically the file system updates the directory information which removes that file. Then the file system notes that the blocks which that file previously occupied are now freed. Subsequently, when you query the file system about how much free space is available, the space occupied by the deleted files appears as part of the free space, since the file system knows it can overwrite that space.

However, the file system does not inform the block device underneath (the SAN/iQ volume) that there is freed-up space. In fact, no mechanism exists to transmit that information. There is no SCSI command which says “Block 198646 can be safely forgotten.” At the block device level, there are only reads and writes.

So, to ensure that our iSCSI block devices work correctly with file systems, any time a block is written to, that block is forever marked as allocated. The file system reviews its “available blocks” list and reuses blocks that have been freed. Consequently, the file system view (such as Windows Disk Management) may show you have X amount of free space, and the CMC view may show the Used Space as 100% used.

-
- △ **CAUTION:** Some file systems support defragmenting which essentially reorders the data on the block device. This can result in the SAN allocating new storage to the volume unnecessarily. Therefore, do not defragment a file system on the SAN unless the file system requires it.
-

Changing the volume size on the server

- △ **CAUTION:** Decreasing the volume size is not recommended. If you shrink the volume in the CMC before shrinking it from the server file system, your data will be corrupted or lost.
-

When you increase the size of the volume on the SAN, you must also increase the corresponding volume, or LUN, on the server side.

Increasing the volume size in Microsoft Windows

After you have increased the volume size on the SAN, you must next expand the Windows partition to use the full space available on the disk.

Windows Logical Disk Manager, the default disk management program that is included in any Windows installation, uses a tool called Diskpart.exe to grow volumes from within Windows. Diskpart.exe is an interactive command line executable which allows administrators to select and manipulate disks and partitions. This executable and its corresponding documentation can be downloaded from Microsoft if necessary.

Follow the steps below to extend the volume you just increased in the SAN.

1. Launch Windows Logical Disk Manager to rescan the disk and present the new volume size.
2. Open a Windows command line and run diskpart.exe.
3. List the volumes that appear to this host by typing the command `list volume`.
4. Select the volume to extend by typing `select volume #` (where # is the corresponding number of the volume in the list).
5. Enter `extend` to grow the volume to the size of the full disk that has been expanded.
Notice the asterisk by the volume and the new size of the volume. The disk has been extended and is now ready for use.

All of the above operations are performed while the volumes are online and available to users.

Increasing the volume size in other environments

Environments other than Windows use alternative disk management tools, which use a utility called Extpart.exe. The only major difference is that instead of selecting the volume number, as in Diskpart.exe, you select the drive letter instead.

Changing configuration characteristics to manage space

Options for managing space on the cluster include

- Changing snapshot retention—retaining fewer snapshots requires less space
- Changing schedules to snapshot a volume—taking snapshots less frequently requires less space
- Deleting volumes or moving them to a different cluster

NOTE: Deleting files on a file system does not free up space on the SAN volume. For more information, see [“Block systems and file systems” \(page 151\)](#). For file-level capacity management, use application or file system-level tools.

12 Using volumes

A volume is a logical entity that is made up of storage on one or more storage systems. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server. Create volumes on clusters that contain one or more storage systems.

Before creating volumes, plan your strategies for using the volume: how you plan to use it, its size, how servers will access it, and how you will manage backups of the data, whether through Remote Copy or third-party applications, or both.

Volumes and server access

After you create a volume, assign it to one or more servers to provide access to volumes by application servers. For detailed information, see [“Controlling server access to volumes” \(page 196\)](#).

Prerequisites

Before you create a volume, you must have created a management group and at least one cluster. For more information, see the following:

- [“Working with management groups” \(page 103\)](#)
- [“Creating a cluster” \(page 132\)](#)

Planning volumes

Planning volumes takes into account multiple factors.

- The number of volumes you need
- The type of volumes you are creating - primary or remote
- The size for each volume
- Whether you plan to use snapshots
- The level of data protection required
- Whether you plan to grow the volume, or leave it the same size

NOTE: If you plan to mount file systems, create a volume for each file system you plan to mount. Then grow each file system independently.

Planning how many volumes

For information about the recommended maximum number of volumes and snapshots that can be created in a management group, see [“Configuration Summary overview” \(page 108\)](#).

Planning volume types

- Primary volumes are volumes used for data storage.
- Remote volumes are used as targets for Remote Copy for business continuance, backup and recovery, and data mining/migration configurations. See the *HP P4000 Remote Copy User Guide* for detailed information about remote volumes.
- A SmartClone volume is a type of volume that is created from an existing volume or snapshot. SmartClone volumes are described in [“SmartClone volumes” \(page 177\)](#).

Guide for volumes

When creating a volume, you define the following characteristics.

Table 43 Characteristics for new volumes

Volume characteristic	Configurable for Primary or Remote Volume	What it means
Basic Tab		
Volume Name	Both	The name of the volume that is displayed in the CMC. A volume name is from 1 to 127 characters and is case sensitive. The volume name cannot be changed. You can enable and customize a default naming convention for volumes. See “Setting naming conventions” (page 15) for more information.
Description	Both	[Optional] A description of the volume.
Size	Primary	<p>The logical block storage size of the volume. Hosts and file systems operate as if storage space equal to the volume size is available in the cluster. This volume size may exceed the true allocated disk space on the cluster for data storage, which facilitates adding more storage systems to the cluster later for seamless storage growth. However, if the volume size does exceed true allocated disk space, the ability to make snapshots may be impacted. See “Using snapshots” (page 161).</p> <p>Remote volumes contain no data, since they serve as pointers to tell the system where to make a copy of a primary snapshot. Therefore, remote volumes do not have a size.</p>
Servers	Both	[Optional] Servers are set up in the management group to connect application hosts to volumes. Select the server that you want to have access to the volume you are creating.
Advanced Tab		
Cluster	Both	If the management group contains more than one cluster, you must specify the cluster on which the volume resides.
Data Protection Level	Both	<p>The data protection level indicates the number and configuration of data copies created on storage systems in the cluster.</p> <p>There are six levels of data protection</p> <ul style="list-style-type: none"> • Network RAID-0 (None) • Network RAID-5 (Single Parity) • Network RAID-6 (Dual Parity) • Network RAID-10 (2-Way Mirror) • Network RAID-10+1 (3-Way Mirror) • Network RAID-10+2 (4-Way Mirror) <p>The default value = Network RAID-10. For information about the data protection levels, see “Planning data protection” (page 141).</p>

Table 43 Characteristics for new volumes *(continued)*

Volume characteristic	Configurable for Primary or Remote Volume	What it means
Type	Both	<ul style="list-style-type: none"> Primary volumes are used for data storage. Remote volumes are used for configuring Remote Copy for business continuance, backup and recovery, or data mining/migration. <p>Default value = Primary</p>
Provisioning	Primary	<ul style="list-style-type: none"> Fully provisioned volumes are the same size on the SAN as the size presented to the application server. Thinly provisioned volumes have less space reserved on the SAN than the size presented to the application server. As data is stored on the volume, the SAN/iQ software automatically increases the amount of space allocated on the SAN. <p>The default value = Full for the following data protection levels.</p> <ul style="list-style-type: none"> Network RAID-0 (None) Network RAID-10 (2-Way Mirror) Network RAID-10+1 (3-Way Mirror) Network RAID-10+2 (4-Way Mirror) <p>The default value = Thin for the following data protection levels.</p> <ul style="list-style-type: none"> Network RAID-5 (Single Parity) Network RAID-6 (Dual Parity) <p>Thin provisioning is the best practice configuration for Network RAID-5 and Network RAID-6 volumes.</p> <p>NOTE: The SAN/iQ software allocates space as needed. However, thin provisioning carries the risk that, if all warnings are ignored, an application server will fail a write because the SAN has run out of disk space.</p>

Creating a volume

A volume resides on the storage systems contained in a cluster. You can easily create a basic volume, or customize the Advanced settings. Both options are described in the following steps.

1. Log in to the management group in which you want to create a volume.
2. In the navigation window, select the cluster in which you want to create a volume.
3. Click **Cluster Tasks**, and select **New Volume**.

Creating a basic volume

You can create a basic volume simply by entering a name and a size for the volume.

1. Enter a name for the volume.
2. [Optional] Enter a description of the volume.
3. Designate a size for the volume.
4. [Optional] Assign a server to the volume.
5. Click OK.

The SAN/iQ software creates the volume. The volume is selected in the navigation window and the Volume tab view displays the Details tab.

NOTE: The system automatically factors data protection levels into the settings. For example, if you create a fully provisioned 500 GB volume and the data protection level is Network RAID-10 (2-Way Mirror), the system automatically allocates 1000 GB for the volume.

To set advanced characteristics for a volume, continue on the Advanced tab of the New Volume window.

Configuring advanced volume settings [optional]

Set additional characteristics for a volume in the Advanced tab in the New Volume window. Advanced settings include the following:

- Cluster (changing the cluster is typically used to migrate volumes to a different cluster at some later time)
- Data protection level
- Volume type
- Provisioning

Descriptions of these characteristics are found in [“Characteristics for new volumes” \(page 155\)](#).

Configuring advanced volume settings

Configure the Advanced settings when you create the new volume if you do not want to use the default settings.

1. Click the Advanced tab on the New Volume window.
2. Change the desired characteristics and click OK when you are finished.

Configuring a volume for either Network RAID-5 or Network RAID-6 automatically creates a snapshot schedule. This default schedule is named *VolumeName_Schedule*, recurs once daily, and retains 1 copy. The first snapshot is created one day after the schedule is created. You can edit this schedule as desired, as described in [“Editing scheduled snapshots” \(page 167\)](#).

Volumes map view

After you create volumes and finish setting up the SAN, use the Map View tab for viewing the relationships between volumes, snapshots, servers, and remote copies. For more information on using the map view tools, see [“Using the display tools” \(page 14\)](#).

Volumes map views include the following:

- Volume Hierarchy
- Remote Copies
- Remote Copies (Volume Hierarchies)
- Servers to Volumes

Editing a volume

When editing a primary volume, you can change the description, size, and advanced characteristics such as the cluster, data protection level, type and provisioning.

NOTE: Moving the volume to a different cluster requires restriping the data in both clusters. Restriping can take hours, or even days.

Table 44 Requirements for changing volume characteristics

Item	Requirements for Changing
Description	Must be from 1 to 127 characters.
Server	Server must have already been created in the management group.
Cluster	<p>The target cluster must</p> <ul style="list-style-type: none"> • Reside in the same management group. • Have sufficient storage systems and unallocated space for the size and data protection level of the volume being moved. • Use a Virtual IP if the originating cluster has a Virtual IP <p>The volume resides on both clusters until all of the data is moved to the new cluster. This causes a restripe of the data on both clusters. For example, you restructure your storage and create an additional cluster. You want to migrate an existing volume to the new cluster as part of the restructuring.</p> <p>NOTE: The volume remains fault-tolerant while being moved.</p>
Data protection level	The cluster must have sufficient storage systems and unallocated space to support the new data protection level. For example, you just added more storage to a cluster and have more capacity. You decide to change the data protection level for a volume from Network RAID-0 to Network RAID-10 to ensure you have redundancy for your data.
Size	<p>Before changing the size of a volume, refer to “Changing the volume size on the server” (page 152)</p> <p>To increase the size of the volume:</p> <ul style="list-style-type: none"> • If you have enough free space in the cluster, simply enter the new size • If you do not have enough free space in the cluster, delete volumes and/or snapshots, or add a storage system to the cluster <p>To decrease the size of the volume (not recommended)</p> <ul style="list-style-type: none"> • If the volume has been or is mounted by any operating system, you must shrink the file system on the volume before shrinking the volume in the CMC. • You also should not decrease the size of the volume below the size needed for data currently stored on the volume.

⚠ CAUTION: Decreasing the volume size is not recommended. If you shrink the volume in the CMC before shrinking it from the server file system, your data will be corrupted or lost.

To edit a volume

1. In the navigation window, select the volume you want to edit.
2. Click **Volume Tasks** and select **Edit Volume**.

Changing the volume description

1. In the Description field, edit the description.
2. Click **OK** when you are finished.

Changing the cluster

If you change the cluster when editing a volume, the change causes the volume and all its data to move to the target cluster.

Requirement

Either before or after changing the cluster, you must stop any applications that are accessing the volume and log off all associated iSCSI sessions.

Even if using the HP DSM for MPIO, log off the volumes from the server, add the VIP or the individual IP addresses of the storage systems in the other cluster, discover and mount volumes.

1. On the Edit Volume window, select the **Advanced** tab.
2. In the Cluster drop-down list, select a different cluster.
3. Click **OK**.

Changing the data protection level

1. In the Data Protection Level drop-down list, select the level of Network RAID you want.

Changing a volume to either Network RAID-5 or Network RAID-6 requires that volume to have a snapshot schedule. If the volume does not already have a snapshot schedule related to it, one will be created automatically. This default schedule is named *VolumeName_Schedule*, recurs once daily, and retains 1 copy. You can edit this schedule as desired, as described in [“Editing scheduled snapshots” \(page 167\)](#).

2. Click OK when you are finished.

Changing the size

1. In the Size field, change the number and change the units if necessary.
2. Click OK when you are finished.

Δ CAUTION: Decreasing the volume size is not recommended. If you shrink the volume in the CMC before shrinking it from the server file system, your data will be corrupted or lost.

Deleting a volume

Delete a volume to remove that volume’s data from the storage system and make that space available. Deleting a volume also deletes all the snapshots underneath that volume, except for clone points and shared snapshots. For more information, see [“Clone point” \(page 184\)](#) and [“Shared snapshot ” \(page 186\)](#).

Δ CAUTION: Deleting a volume permanently removes that volume’s data from the storage system.

Restrictions on deleting volumes

You cannot delete a volume when the volume has a schedule that creates remote copies. You must delete the remote copy schedule first.

Δ CAUTION: Typically, you do not want to delete individual volumes that are part of a volume set. For example, you may set up Exchange to use two volumes to support a StorageGroup: one for mailbox data and one for logs. Those two volumes make a volume set. Typically, you want keep or delete all volumes in a volume set.

Prerequisites

Stop any applications that are accessing the volume and log off all associated iSCSI sessions.

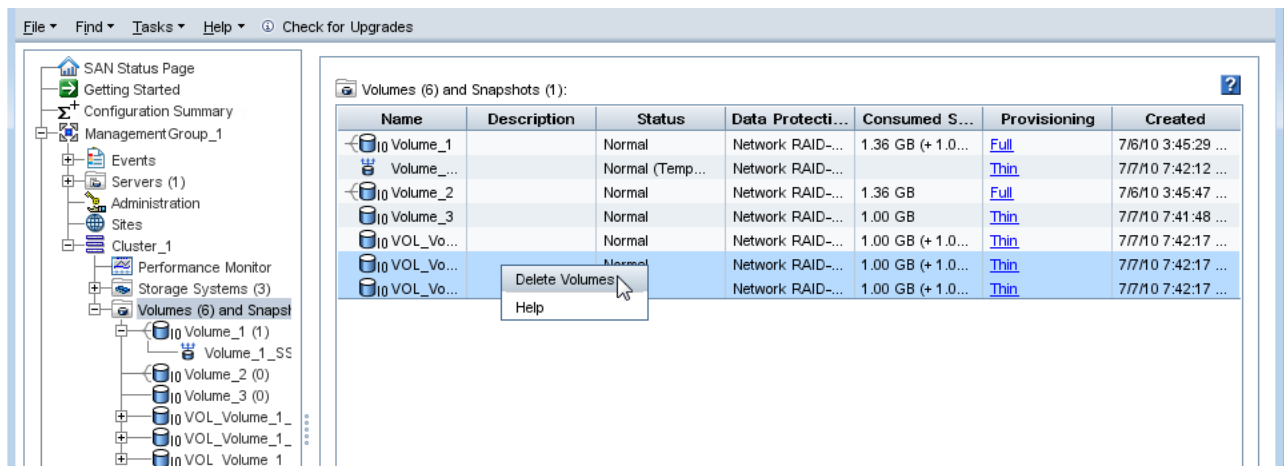
To delete the volume

1. In the navigation window, select the volume you want to delete.
2. Click **Volume Tasks** and select **Delete Volume**.
3. Click **OK**.

To delete multiple volumes

1. In the navigation window, select Volumes and Snapshots.
2. Select the volumes and snapshots you want to delete.

Figure 66 Deleting multiple volumes in one operation



3. Right-click and select **Delete Volumes**.

A warning message opens, asking you to verify that you want to delete the volumes and all the data on them.



4. Select the check box to confirm the deletion and click Delete.
5. The volumes, their associated snapshots (except for clone points and shared snapshots) are deleted from the cluster.

13 Using snapshots

Snapshots are a copy of a volume for use with backup and other applications.

Types of snapshots

Snapshots are one of the following types:

- Regular or point-in-time  — Snapshot that is taken at a specific point in time. However, an application writing to that volume may not be quiesced. Thus, data may be in flight or cached and the actual data on the volume may not be consistent with the application's view of the data.
- Application-managed  — Snapshot of a volume that is taken while the application that is serving that volume is quiesced. Because the application is quiesced, the data in the snapshot is consistent with the application's view of the data. That is, no data was in flight or cached waiting to be written. This type requires the use of the HP P4000 Application Aware Snapshot Manager. For more information, see [“Prerequisites for application-managed snapshots” \(page 163\)](#).

When an application uses two or more volumes, those associated volumes are called a volume set. For example, you may set up Exchange to use two volumes to support a StorageGroup: one for mailbox data and one for logs. Those two volumes make a volume set.

Snapshots versus backups

Backups are typically stored on different physical devices, such as tapes. Snapshots are stored in the same cluster as the volume. Therefore, snapshots protect against data deletion, but not device or storage media failure. Use snapshots along with backups to improve your overall data backup strategy.

Uses and best practices for snapshots

You create snapshots from a volume on the cluster. At any time you can roll a volume back to a specific snapshot, create a SmartClone volume, or use Remote Copy from a snapshot. You can mount a snapshot to a different server and recover data from the snapshot to that server. When planning to use snapshots, ensure that there is space available on the cluster to create the snapshot. For information about snapshots and cluster space, see [“The effect of snapshots on cluster space” \(page 146\)](#)

Snapshots can be used for these cases:

- Source for creating backups
Best Practice—Plan to use a single snapshot and delete it when you are finished.
- Data or file system preservation before upgrading software
Best Practice—Plan to use a single snapshot and delete it when you are finished.
- Protection against data deletion
Best Practice—Plan to use a series of snapshots, deleting the oldest on a scheduled basis. When planning this schedule, consider the following questions:
 - Is space available on the cluster to create the snapshots?
 - What is the optimum schedule and retention policy for this schedule to snapshot a volume? See [“Planning snapshots” \(page 162\)](#) for the average daily change rates for some common applications.For example, if you are using these backups as part of a disaster recovery plan, you might schedule a daily snapshot of the volume and retain 7 copies. A second schedule

would run weekly and retain 5 copies. A third schedule would run monthly and keep 4 copies.

- File-level restore without tape or backup software
 - Source volumes for data mining, test and development, and other data use.
- Best Practice—Use SmartClone volumes. See [“SmartClone volumes” \(page 177\)](#) .

Planning snapshots

When planning to use snapshots, consider their purpose and size. If you are planning to use schedules to snapshot a volume, see [“Clusters and storage systems” \(page 132\)](#) and [Table 45 \(page 162\)](#) for approximate data change rates for some common applications.

Table 45 Common applications’ daily change rates

Application	Daily Change Rates
Fileshare	1 - 3%
Email/Exchange	10 - 20%
Database	10%

NOTE: When considering the size of snapshots in the cluster, remember that the data protection level of the volume is duplicated in the snapshot.

Also consider how many volumes and snapshots you want to create and manage. For information about the recommended maximum number of volumes and snapshots that can be created in a management group, see [“Configuration Summary overview” \(page 108\)](#) and [“Working with management groups” \(page 103\)](#).

Configuring snapshots

When creating a snapshot, you define the following characteristics or options.

Table 46 Snapshot characteristics

Snapshot parameter	What it means
Application-Managed Snapshot	This option quiesces applications on the server before the SAN/iQ software creates the snapshot. This option requires the use of the Application Aware Snapshot Manager. For more information, see “Prerequisites for application-managed snapshots” (page 163) . If the Application Aware Snapshot Manager is not installed, the SAN/iQ software creates a point-in-time snapshot (not using the Application Aware Snapshot Manager).
Snapshot Name	The name of the snapshot that is displayed in the CMC. A snapshot name must be from 1 to 127 characters and is case sensitive. Snapshots have a default naming convention enabled when the CMC is installed. You can change or disable this naming convention. See “Setting naming conventions” (page 15) for information about this naming convention. The following are illegal characters: , ' " ; : =.
Description	(Optional) A description of the snapshot – from 0 to 127 characters.
Assign and Unassign Servers	(Optional for Windows) Configure server access to the snapshot. Required for VMware: In the CMC, first configure a server with the controlling IP address. The controlling IP address is the IP address of the application server on which the

Table 46 Snapshot characteristics *(continued)*

Snapshot parameter	What it means
	vCenter Server is installed. See the <i>HP P4000 Application Aware Snapshot Manager Deployment Guide</i> for more information about the controlling server IP address.

Prerequisites for application-managed snapshots

Creating an application-managed snapshot using the SAN/iQ software is the same as creating any other snapshot. However, you must select the **Application-Managed Snapshot** option in the New Snapshot window. You can create application-managed snapshots for both single and scheduled snapshots. See the *HP P4000 Application Integration Solution Pack Deployment Guide* for server-side requirements for installing and configuring the Application Aware Snapshot Manager.

The following are required for application-managed snapshots:

Table 47 Prerequisites for application-managed snapshots

All	Windows	VMware
<ul style="list-style-type: none"> CMC or CLI latest update HP P4000 Application Integration Solution Pack, specifically the HP P4000 Application Aware Snapshot Manager (latest update) installed on the application server (See the <i>HP P4000 Application Integration Solution Pack User Guide</i>) Management group authentication set up for the Application Aware Snapshot Manager (see the <i>HP P4000 Application Integration Solution Pack User Guide</i>) Application on the server that is can be quiesced Server configured with iSCSI connection (see “Controlling server access to volumes” (page 196)) 	<ul style="list-style-type: none"> SAN/iQ software 8.5 or later for Windows application-managed snapshots Microsoft iSCSI initiator 	<ul style="list-style-type: none"> SAN/iQ software 9.5 or later for VMware application-managed snapshots CMC Server configured with iSCSI connection that must use the IP address of the vCenter Server as the Controlling Server IP address ESX Server software iSCSI initiator

Application-managed snapshots for volume sets

When you create an application-managed snapshot of a volume in a volume set, the software recognizes that the volume is part of a volume set, and prompts you to create a snapshot for each volume in the volume set. The result is a snapshot set that corresponds to the volume set. To see any associated snapshots, select a snapshot, click the Details tab, and look at the Snapshot Set field.

NOTE: After you create snapshots for a volume set, typically you do not want to delete individual snapshots from the snapshot set. You want to keep or delete all snapshots for the volume set. If you need to roll back to a snapshot, typically you want to roll back each volume in the volume set to its corresponding snapshot. The system gives you the option to automatically delete or roll back all associated volumes.

Creating snapshots

Create a snapshot to preserve a version of a volume at a specific point in time. For information about snapshot characteristics, see “Configuring snapshots” (page 162).

Creating an application-managed snapshot, with or without volume sets, requires the use of the Application Aware Snapshot Manager. The application-managed snapshot option quiesces Windows and VMware applications on the server before creating the snapshot. For more information, see [“Prerequisites for application-managed snapshots” \(page 163\)](#). If the Application Aware Snapshot Manager is not installed, the SAN/iQ software creates a regular, point-in-time snapshot.

Creating regular or application-managed snapshots

The snapshot creation process for application-managed snapshots differs when a Windows application has associated volumes. See [“Creating snapshots for volume sets” \(page 164\)](#). When using application-managed snapshots with VMware vCenter Server, you must first create a server in the CMC and use the IP address of the application server on which the vCenter Server is installed as the controlling server IP address in the New Server window.

1. Log in to the management group that contains the volume for which you want to create a new snapshot.
2. Right-click the volume, and select **New Snapshot**.
3. (Optional) If you are using the Application Aware Snapshot Manager and want to quiesce the application before creating the snapshot, select **Application-Managed Snapshot**.
The system fills in the Description field and disables the Assign and Unassign Servers button automatically. You can assign servers after the snapshot is created.
4. Enter a name for the snapshot, or accept the default.
5. (Optional) Enter a description of the snapshot.
6. (Optional) Assign a server to the snapshot.
7. Click **OK** when you are finished.

NOTE: In the navigation window, snapshots are listed below the volume in descending date order, from newest to oldest.

Creating snapshots for volume sets

1. Log in to the management group that contains the volume for which you want to create a new snapshot.
2. Right-click the volume, and select **New Snapshot**.
3. Select **Application-Managed Snapshot**.
The software fills in the Description field and disables the Assign and Unassign Servers button automatically. You can assign servers after the snapshot is created.
4. Enter a name for the snapshot or accept the default.
5. Click **OK**.
The New Snapshot – Associated Volumes window opens with a list of all volumes in the volume set.
6. (Optional) Edit the Snapshot Name and Description for each snapshot.

NOTE: Be sure to leave the Application-Managed Snapshots check box selected. This option maintains the association of the volumes and snapshots and quiesces the application before creating the snapshots. If you clear the check box, the system creates a point-in-time snapshot of each volume listed.

7. Click **Create Snapshots** to create a snapshot of each volume.
The snapshots appear in the CMC. To see the associated snapshots, select a snapshot, click the **Details** tab, and look at the Snapshot Set field.

Editing a snapshot

You can edit both the description of a snapshot and its server assignment. The description must be from 0 to 127 characters.

1. Log in to the management group that contains the snapshot that you want to edit.
2. In the navigation window, select the snapshot.
3. Click **Snapshot Tasks** on the Details tab, and select **Edit Snapshot**.
4. Change the description as necessary.
5. Change the server assignment as necessary.
6. Click **OK** when you are finished.

Scheduling snapshots

Use schedules to create a series of snapshots up to a specified number, or for a specified time period. After the specified time period, or accumulated number of snapshots, the earliest snapshot is deleted when the new one is created. For example, you plan to keep a series of daily snapshots for one week, up to five snapshots. After creating the sixth snapshot, the earliest snapshot is deleted, thereby keeping the number of snapshots on the volume at five. A schedule can also create a single snapshot and then delete it when it is no longer needed. Snapshot schedules can also be paused and resumed.

Scripting snapshots, either recurring or single snapshots, can also take place on the server side. Scripted snapshots offer greater flexibility for quiescing hosts while taking snapshots, and for automating tasks associated with volumes and their snapshots.

Best practices for scheduling snapshots of volumes

- If you do not have an NTP server configured, before you create the schedule, you should refresh the time setting of the management group to ensure that the storage systems are all set to the correct time.
- Configure schedules to snapshot a volume during off-peak hours. If setting schedules for multiple volumes, stagger the schedules with at least an hour between start times for best results.

Table 48 Planning the scheduling for snapshots

Requirement	What it means
Plan for capacity management	<p>Scheduling snapshots should be planned with careful consideration for capacity management as described in “Managing capacity using volume size and snapshots” (page 147). Pay attention to how you want to retain snapshots and the capacity in the cluster. If you want to retain <n> snapshots, the cluster should have space for <n+1>.</p> <p>It is possible for the new snapshot and the one to be deleted to coexist in the cluster for some period of time.</p> <p>If there is not sufficient room in the cluster for both snapshots, the scheduled snapshot will not be created, and the snapshot schedule will not continue until an existing snapshot is deleted or space is otherwise made available.</p>
Plan scheduling and retention policies	<p>The minimum recurrence you can set for snapshots is 30 minutes. The maximum number of snapshots (scheduled and manual combined) you can retain is 50 snapshots per volume. There are practical limits to the number of snapshots that a particular SAN can support and still maintain adequate performance. For information on optimum configuration limits, performance, and scalability, see “Configuration Summary overview” (page 108).</p>

Requirements for snapshot schedules

Table 49 Characteristics for creating a schedule to snapshot a volume

Item	Description and requirements
Name	<p>The name of the schedule that is displayed in the CMC. A scheduled snapshot name must be from 1 to 127 characters and is case sensitive. Snapshots created by a schedule have a default naming convention enabled when the CMC is installed. You can change or disable this naming convention. See “Setting naming conventions” (page 15) for information about this naming convention.</p> <p>The name you enter in the Create Schedule to Snapshot a Volume window will be used with sequential numbering. For example, if the name is Backup, the list of snapshots created by this schedule will be named Backup.1, Backup.2, Backup.3.</p>
Description	[Optional] Must be from 0 to 127 characters.
Start at	The start date and time is usually set for the future. However, it can be set to occur in the past. If set for a past time and date, a snapshot is created immediately with the parameters set in the schedule.
Recurrence	The recurrence can be set to every <i>n</i> number of minutes, hours, days, or weeks, or to never recur. The minimum recurrence is 30 minutes.
Application-managed snapshot	[Optional] This option quiesces applications on the application server before the SAN/iQ software creates the snapshot.
Retention	The retention criteria can be for a specified number of snapshots, up to 50, or for a designated period of time.

Scheduling snapshots for volume sets

When you create a schedule to snapshot a volume that has associated volumes, the system automatically creates snapshots for each associated volume. For information about volume sets, see [“Application-managed snapshots for volume sets” \(page 163\)](#).

The schedule also reflects the volume associations based on the volume you select when you create the schedule. That volume becomes the “owning” volume. The Volume Set field of the schedule displays (O) next to the owning volume. You should check that the field displays all of the volumes that you want to snapshot. It is possible that the owning volume is not aware of all associated volumes. If it is not, select a volume that is aware of all associated volumes, and create the schedule there.

Updating schedule for volume sets

When you first create the schedule, the system stores information about the volume set as it exists at that time. If you add volumes to or remove volumes from the volume set using the application, you must update the schedule. To update it, you only need to edit the schedule and click **OK**. The system automatically updates the volume set information when you click **OK**. If you want to see the updated information, you can click **Verify Volume Associations**, then click **OK**. For more information, see [“Editing scheduled snapshots” \(page 167\)](#).

NOTE: If you have a schedule to **remote** snapshot a volume and you add a volume to the volume set using the application, the system cannot update the volume set information as described above. You must delete the schedule and create a new one to reflect the current volume set.

Creating a schedule to snapshot a volume

1. In the navigation window, select volume for which you want to create a schedule for snapshots.
2. Click **Volume Tasks** on the Details tab and select **New Schedule to Snapshot a Volume**.
3. Enter a name for the schedule.
4. (Optional) Enter a snapshot description.
5. Click **Edit** to specify a start date and time.

The Date and Time Configuration window opens. Use this window to set the date and time for the first snapshot created by this schedule.

6. Click **OK** when you are finished setting the date and time.
7. Select a recurrence schedule.
8. If you want to quiesce the application before creating the snapshot, select **Application-Managed Snapshot**.

This option requires the use of the Application Aware Snapshot Manager. For more information, see [“Prerequisites for application-managed snapshots” \(page 163\)](#). If the Application Aware Snapshot Manager is not installed, the SAN/iQ software creates a point-in-time snapshot.

9. Specify the retention criteria for the snapshot.
10. Click **OK** when you have finished creating the schedule.

If the volume is <i>not</i> part of a volume set	If the volume is part of a volume set
To view the schedule, select the Schedules tab view.	<p>The Volume Associations Found window opens. This window informs you that the volume you are creating a schedule for has one or more associated volumes. The system will create snapshots for each associated volume. For information about volume sets, see “Application-managed snapshots for volume sets” (page 163).</p> <p>Click Create Schedule to complete the process and create scheduled snapshots for each volume.</p> <p>When you edit the schedule, the Volume Set field lists the volumes that are part of the set, and displays (O) next to the volume that owns the schedule. For more information, see “Scheduling snapshots for volume sets” (page 166).</p> <p>If you do not want to create scheduled snapshots for each volume in the volume set, click Cancel.</p> <p>NOTE: If the volumes associated with this schedule change (add or remove volumes), you can update the volume information by editing the schedule. For more information, see “Editing scheduled snapshots” (page 167)</p>

Editing scheduled snapshots

You can edit everything in the scheduled snapshot window except the name.

If the snapshot is part of a snapshot set, you can also verify that the volumes included in the schedule are the current volumes in the volume set. For more information, see [“Scheduling snapshots for volume sets” \(page 166\)](#).

1. In the navigation window, select the volume for which you want to edit the scheduled snapshot.
2. In the tab window, click the **Schedules** tab to bring it to the front.
3. Select the schedule you want to edit.
4. Click **Schedule Tasks** on the Details tab, and select **Edit Schedule**.
5. Change the desired information.

6. (Optional) If displayed, click **Verify Volume Associations** to see if the volume set included in the snapshot set is up to date.

The Verify Volume Associations window opens, showing the volumes currently associated with the schedule. Any volumes that have been added to or removed from the volume set are reflected. Click **Close** to return to the Edit Schedule to a Snapshot Volume window. The updated list of volumes is populated in the Volume Set field. For more information, see [“Scheduling snapshots for volume sets” \(page 166\)](#).

This lets you see the current volume set information. The information is automatically updated when you click **OK**.

7. Click **OK**.

Pausing and resuming scheduled snapshots

At times it may be convenient to prevent a scheduled snapshot from taking place. When you pause a snapshot schedule, the snapshot deletions for that schedule are paused as well. When you resume the schedule, both the snapshots and the snapshot deletions resume according to the schedule.

Pause a schedule

1. In the navigation window, select the volume for which you want to pause the snapshot schedule.
2. Click the **Schedules** tab to bring it to the front.
3. Select the schedule you want.
4. Click **Schedule Tasks** on the Details tab, and select **Pause Schedule**.
5. In the Confirm window, click **OK**.

In the Next Occurrence column of the Schedules tab window, this snapshot schedule is marked as paused.

6. Make a note to resume this snapshot schedule at a convenient time.

Resume a schedule

1. In the navigation window, select the volume for which you want to resume the snapshot schedule.
2. Click the **Schedules** tab to bring it to the front.
3. Select the schedule you want.
4. Click **Schedule Tasks** on the Details tab, and select **Resume Snapshot Schedule**.
5. In the Confirm window, click **OK**.

In the Next Occurrence column of the tab window, this snapshot schedule shows the date and time the next snapshot will be created.

Deleting schedules to snapshot a volume

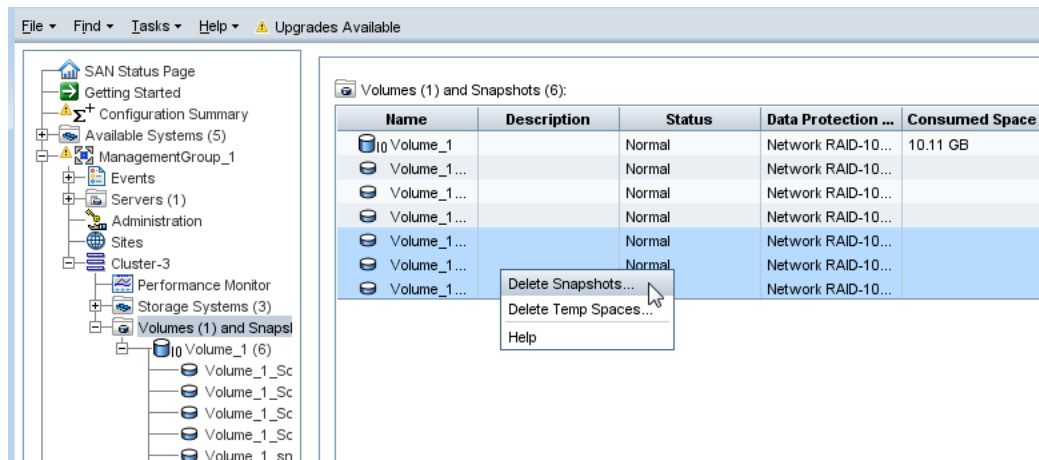
NOTE: After you delete a snapshot schedule, if you want to delete snapshots created by that schedule, you must do so manually.

1. In the navigation window, select the volume for which you want to delete the snapshot schedule.
2. Click the **Schedules** tab to bring it to the front.
3. Select the schedule you want to delete.
4. Click **Schedule Tasks** on the Details tab, and select **Delete Schedule**.
5. To confirm the deletion, click **OK**.

The Schedules tab refreshes without the deleted snapshot schedule.

6. [Optional] To delete snapshots related to that schedule, select the **Volumes and Snapshots** node where you can delete multiple snapshots from a list.

Figure 67 Delete multiple snapshots from the volumes and snapshots node



Scripting snapshots

Application-based scripting allows automatic snapshots of a volume. For detailed information, see [“Working with scripting” \(page 195\)](#) and the *HP CLIQ - The SAN/iQ Command-Line Interface User Guide*, for information about the SAN/iQ software command-line interface.

Mounting a snapshot

A snapshot is a copy of a volume. To access data in the snapshot, you have two choices:

- Create a SmartClone volume from the snapshot to use for data mining, development and testing, or creating multiple copies. See [“Create a new SmartClone volume from the snapshot” \(page 174\)](#).
- Mount the snapshot for backing up or data recovery. You assign the snapshot to a server as a read/write volume and connect to it with an iSCSI initiator.

Mounting a snapshot on a server adds temporary space to the snapshot. See [“Managing snapshot temporary space” \(page 172\)](#) for more detailed information about temporary space.

Mounting the snapshot on a host

You can add a server to the snapshot when it is created, or add the server later.

1. If it is not already added, add the server on which you want to mount the snapshot to the management group.
2. Assign the snapshot to the server, and configure the snapshot for read/write access.
3. Configure server access to the snapshot
4. If you mount a Windows application-managed snapshot as a volume, use `diskpart.exe` to change the resulting volume's attributes, as described in [“Making a Windows application-managed snapshot available” \(page 170\)](#).

When you have mounted the snapshot on a host, you can do the following:

- Recover individual files or folders and restore to an alternate location
- Use the data for creating backups

Making a Windows application-managed snapshot available

If you do any of the following using a Windows application-managed snapshot, you must use `diskpart.exe` to make the resulting volume available:

- Convert temporary space
- Create a SmartClone
- Promote a remote volume to a primary volume
 - Failover/Failback Volume Wizard and selecting the “Failover the Primary Volume to the Selected Remote Volume Below” option
 - Edit Volume and changing a remote snapshot to a primary volume

Making a Windows application-managed snapshot available on a stand-alone server

Use this procedure to make a Windows application-managed snapshot available on a stand-alone server (not part of a Microsoft cluster).

1. Disconnect the iSCSI sessions.
2. Do one of the following (based on what you want to do with the application-managed snapshot):
 - Convert temporary space.
 - Create a SmartClone.
 - Promote a remote volume to a primary volume using:
 - Failover/Failback Volume Wizard and selecting the “Failover the Primary Volume to the Selected Remote Volume Below” option.
 - Edit Volume and changing a remote snapshot to a primary volume.
3. Connect the iSCSI sessions to the new target volume.
4. Launch Windows Logical Disk Manager.
5. Bring the disk online.
6. Open a Windows command line and run `diskpart.exe`.
7. List the disks that appear to this server by typing the command `list disk`.
8. Select the disk you are working with by typing `select disk #` (where # is the corresponding number of the disk in the list).
9. Display the options set at the disk level by typing `detail disk`.
If the disk is listed as read-only, change it by typing `att disk clear readonly`.
10. Select the volume you are working with by typing `select volume #` (where # is the corresponding number of the volume in the list).
11. Display the volume's attributes by typing `att vol`.
The volume will show that it is hidden, read-only, and shadow copy.
12. Change these attributes by typing `att vol clear readonly hidden shadowcopy`.
13. Exit `diskpart` by typing `exit`.
14. Reboot the server.
15. Verify that the disk is available by launching Windows Logical Disk Manager.
You may need to assign a drive letter, but the disk should be online and available for use.
16. If the server is running Windows 2008 or later and you promoted a remote application-managed snapshot to a primary volume, start the HP P4000 CLI and clear the VSS volume flag by typing `clearvssvolume flags volumename=[drive_letter]` (where [drive_letter] is the corresponding drive letter, such as G:).
17. Reboot the server.

Making a Windows application-managed snapshot available on a server in a Microsoft cluster

Use this procedure to make an application-managed snapshot available on servers that are in a Microsoft cluster.

NOTE: We recommend contacting Customer Support before performing this procedure.

1. Disconnect the iSCSI sessions.
2. Do one of the following (based on what you need to do with the application-managed snapshot):
 - Convert temporary space.
 - Create a SmartClone.
 - Promote a remote volume to a primary volume.
 - Failover/Failback Volume Wizard and selecting the “Failover the Primary Volume to the Selected Remote Volume Below” option.
 - Edit Volume and changing a remote snapshot to a primary volume.
3. Connect the iSCSI sessions to the new target volume.
4. Launch Windows Logical Disk Manager.
5. Bring the disk online.
6. Open the system event log and find the IDs for the disks you are working with.

The disks will have new disk IDs. The log will show errors for the disks, along with the IDs the cluster was expecting to see for each disk.
7. Open a Windows command line and run `diskpart.exe`.
8. List the disks that appear to this server by typing the command `list disk`.
9. Select the disk you are working with by typing `select disk #` (where # is the corresponding number of the disk in the list).
10. Display the options set at the disk level by typing `detail disk`.

If the disk is listed as read-only, change it by typing `att disk clear readonly`.

The details show the expected ID for each disk. If the server is running Windows 2003, see Microsoft KB 280425 for how to change the disk IDs.
11. On Windows 2008 and later, change the disk ID to the expected ID by typing `uniqueid disk ID=[expected_ID]` (where [expected_ID] is the corresponding number of the disk in the list).
12. Select the volume you are working with by typing `select volume #` (where # is the corresponding number of the volume in the list).
13. Display the volume's attributes typing `att vol`.

The volume will show that it is hidden, read-only, and shadow copy.
14. Change these attributes by typing `att vol clear readonly hidden shadowcopy`.
15. Exit `diskpart` by typing `exit`.
16. Reboot the server.
17. Verify that the disk is available by launching Windows Logical Disk Manager.

You may need to assign a drive letter, but the disk should be online and available for use.
18. If the server is running Windows 2008 or later and you promoted a remote application-managed snapshot to a primary volume, start the HP P4000 CLI and clear the VSS volume flag by typing `clearvssvolume flags volumename=[drive_letter]` (where [drive_letter] is the corresponding drive letter, such as G:).
19. Reboot the server.

Managing snapshot temporary space

You can either delete the temporary space to free up space on the cluster, or, if you need data that may have been written to the temporary space, convert that temporary space to a SmartClone volume.

Convert the temporary space to access data

Convert the snapshot temporary space if you have written data to a mounted snapshot and you need to permanently save or access that data. Converting the temporary space creates a SmartClone volume that contains the original snapshot data plus any additional data written after the snapshot was mounted.

Prerequisites

Stop any applications that are accessing the snapshot, and log off all related iSCSI sessions

To convert the temporary space:

1. Right-click the snapshot for which you want to save the additional data.
2. Select **Convert Temporary Space** from the menu.
3. Enter a name for the volume and an optional description.
4. Click **OK**.

The temporary space becomes a volume with the name you assigned. The original snapshot becomes a clone point under the new volume. For more information about clone points, see [“Rolling back a volume to a snapshot or clone point” \(page 172\)](#).

5. If you converted temporary space from an application-managed snapshot, use `diskpart.exe` to change the resulting volume's attributes.

For more information, see [“Making a Windows application-managed snapshot available” \(page 170\)](#).

Delete the temporary space

The snapshot temporary space is deleted when the snapshot is deleted. However, you can manually delete the snapshot temporary space if you need to free up space on the cluster. Note that if you have written any data to the snapshot, that data will be deleted along with the temporary space. If you want to save that data, convert the temporary space to a volume.

Prerequisite

Stop any applications that are accessing the snapshot, and log off all related iSCSI sessions.

To delete the temporary space:

1. In the navigation window, select snapshot for which you want to delete the temporary space.
2. Right-click, and select **Delete Temporary Space**.
A warning message opens.
3. Click **OK** to confirm the delete.

Rolling back a volume to a snapshot or clone point

Rolling back a volume to a snapshot or a clone point replaces the original volume with a read/write copy of the selected snapshot. Rolling back a volume to a snapshot deletes any new snapshots that may be present, so you have some options to preserve data in those snapshots.

- Instead of rolling back, use a SmartClone volume to create a new volume from the target snapshot. This volume gets a new name and the target snapshot becomes a clone point, shared

between the original volume and the new SmartClone volume. For detailed information about SmartClone volumes, see [“What are SmartClone volumes?”](#) (page 177).

- Use Remote Copy to copy the newer snapshots that you want to keep, before performing the rollback. See the *HP P4000 Remote Copy User Guide* for more information about copying data.

Best practices for rolling back a volume

- Stop any applications that are accessing the volume, and log off all related iSCSI sessions.
- If a volume is part of a volume set, typically you want to roll back each volume using its corresponding snapshot. The system gives you the option to automatically roll back all associated volumes. To see any associated snapshots, select a snapshot, click the **Details** tab, and look at the Snapshot Set field. For more information, see [“Creating snapshots for volume sets”](#) (page 164).

Restrictions on rolling back a volume

You cannot roll back a volume when a clone point exists that is newer than the snapshot you want to use for rolling back. You can create a SmartClone from the snapshot you want to use, or you must delete all but one volume that depends on that clone point. Once you delete all but one volume that depends on a clone point, the clone point returns to being a standard snapshot.

Rolling back a volume to a snapshot or clone point

You can roll back a specific volume from a clone point. The clone point selected will roll back to the parent volume it is listed under in the navigation view.

⚠ CAUTION: During a rollback, snapshots that are newer than the one you intend to roll back are deleted. You will lose all data stored since the rolled-back snapshot was created. Consider creating a SmartClone volume, or a Remote Copy, before the roll back to preserve that data.

1. Log in to the management group that contains the volume that you want to roll back.
2. In the navigation window, select the snapshot to which you want to roll back.

Review the snapshot Details tab to ensure you have selected the correct snapshot.

3. Click **Snapshot Tasks** on the Details tab, and select **Roll Back Volume**.

A warning message opens that illustrates the possible consequences of performing a rollback, including

- Existing iSCSI sessions present a risk of data inconsistencies.
- All newer snapshots will be deleted.
- Changes to the original volume since the snapshot was created will be lost.

If you do not have connected iSCSI sessions or newer snapshots, those issues will not be reflected in the message.

If the snapshot is <i>not</i> part of a snapshot set	If the snapshot is part of a snapshot set
<p>You have three choices for continuing from this message window:</p> <ul style="list-style-type: none">• Click OK. See “Continue with standard roll back” (page 174).• Click New SmartClone Volume. See “Create a new SmartClone volume from the snapshot” (page 174).• Click Cancel. See “Cancel the rollback operation” (page 175).	<p>You have the following possible choices for continuing from this message window:</p> <ul style="list-style-type: none">• Click Roll Back ALL Associated Volumes. This is recommended. See “Roll back all associated volumes” (page 175).• Click Roll Back Selected Volume Only. See “Continue with standard roll back” (page 174). <p>NOTE: This will only the roll back the volume of the selected snapshot. Click OK. This is not recommended.</p>

If the snapshot is <i>not</i> part of a snapshot set	If the snapshot is part of a snapshot set
	<ul style="list-style-type: none"> Click Roll Back Volume. See “Continue with standard roll back” (page 174). NOTE: This will leave some snapshot sets incomplete. This is not recommended. Click Cancel. See “Cancel the rollback operation” (page 175).

Continue with standard roll back

The following steps result with the original volume, with its original name, returned to the state of the rolled back snapshot. If the snapshot is part of a snapshot set, this is not recommended.

The volume rolls back to the snapshot, deleting any newer snapshots. The rolled back snapshot remains intact underneath the volume and retains the data. Any data that had been added to the volume since the snapshot was created is deleted.

1. If you rolled back an application-managed snapshot, use `diskpart.exe` to change the resulting volume's attributes.

For more information, see [“Making a Windows application-managed snapshot available” \(page 170\)](#).

2. Reconnect iSCSI sessions to the volume, and restart the applications.

Create a new SmartClone volume from the snapshot

Instead of continuing with a standard roll back, you can create a new SmartClone volume, with a new name, from the selected snapshot. This choice preserves any newer snapshots and any new data in the original volume.

1. Click **New SmartClone Volume**.

2. Enter a name, and configure the additional settings.

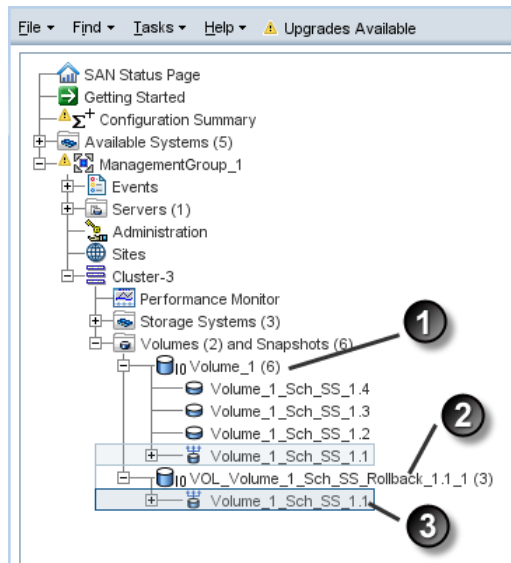
For more information about characteristics of SmartClone volumes, see [“Defining SmartClone volume characteristics” \(page 180\)](#).

3. Click **OK** when you have finished setting up the SmartClone volume and updated the table.

The new volume appears in the navigation window, with the snapshot now a designated clone point for both volumes.

4. Assign a server, and configure hosts to access the new volume, if desired.

Figure 68 New volume with shared clone point



1. Original volume
2. New SmartClone volume from snapshot
3. Shared clone point

5. If you created the SmartClone from an application-managed snapshot, use `diskpart.exe` to change the resulting volume's attributes.

For more information, see [“Making a Windows application-managed snapshot available”](#) (page 170).

Roll back all associated volumes

The recommended method for rolling back a snapshot that is part of a snapshot set is to roll back each volume to its corresponding snapshot. For more information about snapshot sets, see [“Application-managed snapshots for volume sets”](#) (page 163).

1. Click **Roll Back ALL Associated Volumes**.
Each associated volume rolls back to its corresponding snapshot.
2. Use `diskpart.exe` to change the resulting volume's attributes.
For more information, see [“Making a Windows application-managed snapshot available”](#) (page 170).
3. Reconnect iSCSI sessions to the volume, and restart the applications.

Cancel the rollback operation

If you need to log off iSCSI sessions, stop application servers, or other actions, cancel the operation, perform the necessary tasks, and then do the rollback.

1. Click **Cancel**.
2. Perform necessary actions.
3. Start the rollback again.

Deleting a snapshot

When you delete a snapshot, the data necessary to maintain volume consistency are moved up to the next snapshot or to the volume (if it is a primary volume), and the snapshot is removed from the navigation window. The temporary space associated with the snapshot is deleted.

Restrictions on deleting snapshots

You cannot delete a snapshot when the snapshot is:

- A clone point.
- In the process of being deleted or being copied to a remote management group.
- The primary snapshot that is copied using Remote Copy, and you are not logged into the remote management group that it is copied to.

⚠ CAUTION: Typically, you do not want to delete individual snapshots that are part of a snapshot set. To see any associated snapshots, select a snapshot, click the **Details** tab, and look at the Snapshot Set field. For information about snapshot sets, see [“Prerequisites for application-managed snapshots” \(page 163\)](#). Typically, you want to keep or delete all snapshots for a volume set. If you need to roll back to a snapshot, you want to roll back each volume in the volume set to its corresponding snapshot. The system gives you the option to automatically delete or roll back all associated volumes.

CAUTION: Network RAID-5 and Network RAID-6 volumes require snapshots in order to achieve space utilization benefits. This means that deleting the last snapshot of a Network RAID-5 volume causes its space requirement to be the same as a Network RAID-10 (2-Way Mirror) volume. Similarly deleting the last snapshot of a Network RAID-6 volume causes its space requirement to be the same as a Network RAID-10+1 (3-Way Mirror) volume. It is possible, therefore, for the storage cluster not to have enough space to accommodate the snapshot deletion. Deleting the last snapshot of a Network RAID-5 or Network RAID-6 volume is not recommended.

Prerequisites

Stop any applications that are accessing the snapshot, and log off all related iSCSI sessions

To delete the snapshot:

1. Log in to the management group that contains the snapshot that you want to delete.
2. In the navigation window, select the snapshot that you want to delete.
3. Review the Details tab to ensure you have selected the correct snapshot.
4. Click **Snapshots Tasks** on the Details tab, and select **Delete Snapshot**.

If the snapshot is <i>not</i> part of a snapshot set	If the snapshot is part of a snapshot set
A confirmation message opens. Click OK .	A warning message opens. <ul style="list-style-type: none">• To delete all snapshots in the snapshot set, click Delete All Associated Snapshots.• To delete only the snapshot you selected, click Delete Selected Snapshot Only.• To cancel the deletion, click Cancel.

14 SmartClone volumes

SmartClone are space-efficient copies of existing volumes or snapshots. They appear as multiple volumes that share a common snapshot, called a clone point. They share this snapshot data on the SAN. SmartClone volumes can be used to duplicate configurations or environments for widespread use, quickly and without consuming disk space for duplicated data. Use the SmartClone process to create up to 25 volumes in a single operation. Repeat the process to create more volumes, or use the CLI to create larger quantities in a single scripted operation.

What are SmartClone volumes?

SmartClone volumes can be created instantaneously and are fully featured, writable volumes. The only difference between regular volumes, snapshots, and SmartClone volumes is that SmartClone volumes are dependent on the clone point, that is, the snapshot they are created from. Additionally, they may minimize space used on the SAN. For example, you create a volume with a specific OS configuration. Then, using the SmartClone process, you create multiple volumes with access to that same OS configuration, and yet you only need a single instance of the configuration. Only as additional data is written to the different SmartClone volumes do those volumes consume additional space on the SAN. The space you save is reflected on the Use Summary tab in the Cluster tab window, described in [“Cluster use summary” \(page 148\)](#).

Multiple SmartClone volumes can be individually managed just like other volumes. SmartClone volumes can be used long term in production environments. Examples of common uses for SmartClone volumes:

- Deploy large quantities of virtual machine clones, including virtual servers and virtual desktops
- Copy production data for use in test and development environments
- Clone database volumes for data mining
- Create and deploy boot-from-SAN images

Prerequisites

- You must have created a management group, cluster, and at least one volume.
- You must have enough space on the SAN for the configuration you are planning.
- You must be running SAN/iQ software version 8.0 or later.

SmartClone volume terminology

[Table 50 \(page 177\)](#) lists terms and definitions used for the SmartClone volumes feature. The illustration in [Figure 69 \(page 178\)](#) shows how the SmartClone volumes and related elements look in the CMC.

Table 50 Terms used for SmartClone features



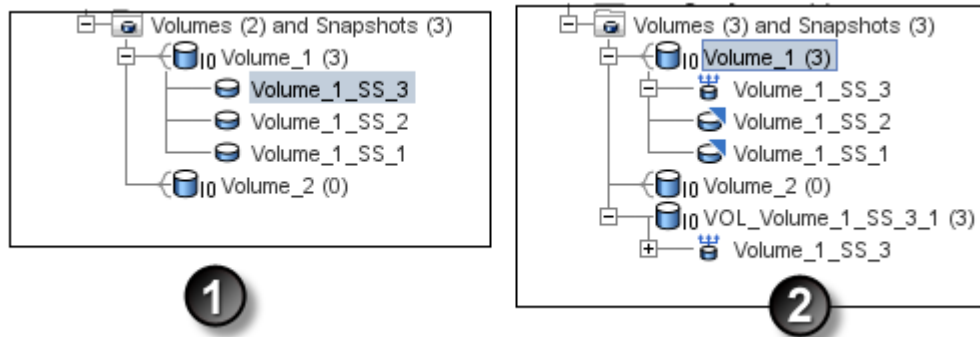
Term	Definition
SmartClone Volume	A volume created using the SmartClone process. In Figure 69 (page 178) , the volume Volume_2 is a SmartClone volume.
Clone point 	The snapshot from which the SmartClone volumes are created. The clone point cannot be deleted. In Figure 69 (page 178) , the snapshot Volume_1_SS_3 is the clone point.
Shared snapshot 	Shared snapshots occur when a clone point is created from a newer snapshot that has older snapshots below it in the tree. Shared snapshots can be deleted. In Figure 69

Table 50 Terms used for SmartClone features *(continued)*

Term	Definition
	(page 178), the snapshots Volume_1_SS_1 and Volume_1_SS_2 are shared snapshots.
Map view	Tab that displays the relationships between clone points and SmartClone volumes. See the map view in Figure 81 (page 190) and Figure 82 (page 191) .

In [Figure 69 \(page 178\)](#) you can see on the left a regular volume with three snapshots and on the right, a regular volume with one SmartClone volume, one clone point, and two shared snapshots.

Figure 69 How SmartClone volumes, clone points, and shared snapshots appear in the CMC



1. Regular volumes and snapshots

2. SmartClone volumes, with clone points and shared snapshots

Example scenarios for using SmartClone volumes

The following examples are just a few of the most typical scenarios for using SmartClone volumes.

Deploy multiple virtual or boot-from-SAN servers

You can save significant space in environments with multiple virtual or boot-from-SAN servers that use the same base operating system. A server's operating system takes up considerable storage but does not change frequently. You can create a master image of the operating system on a volume and prepare it for duplication. Then you can create large quantities of SmartClone volumes from that master image without using additional storage capacity. Each SmartClone volume you create from the master image is a full read/write version of the operating system and has all the same management features as a regular HP P4000 SAN Solution volume.

Scenario: Computer training lab

You run a computer lab for a technical training company. You routinely set up training environments for classes in programming languages, database development, web design, and other applications. The classes are anywhere from 2 days to 1 week long, and your lab can accommodate 75 students.

On your HP P4000 SAN Solution, you maintain master desktop images for each class offering. These desktop images include all the software applications the students need for each class, in the default configuration required for the start of the class.

To prepare for an upcoming class with 50 students, you clone the 50 student desktops from the master image, without consuming additional space on the SAN. You configure the iSCSI connections and the students are ready to start working. During the class, the only additional data added to the SAN is the trainees' class work. When the class is finished, you can roll back all 50 SmartClone volumes to the clone point and recreate the desktops.

Safely use production data for test, development, and data mining

Use SmartClone volumes to safely work with your production environment in a test and development environment, before going live with new applications or upgrades to current applications. Or, clone copies of your production data for data mining and analysis.

Test and development

Using the SmartClone process, you can instantly clone copies of your production LUNs and mount them in another environment. Then you can run new software, install upgrades, and perform other maintenance tasks. When the new software or upgrades testing is complete, either redirect your application to the SmartClone volume you have been using, or delete the SmartClone volume and proceed with the installation or upgrades in the production environment.

Data mining

Assume you want to track monthly trends in web requests for certain types of information. Once a month, you create a SmartClone volume of the web server transaction, mount the volume to a different server, and analyze and track usage or other trends over time. This monthly SmartClone volume takes minimal additional space on the SAN, while providing all the data from the web server database.

Clone a volume

In addition to the cases described above, SmartClone volumes can be created as needed for any purpose. These volumes provide exact copies of existing volumes without the need to provision additional space, until and unless you want to write new data.

Planning SmartClone volumes

Planning SmartClone volumes takes into account multiple factors, such as space requirements, server access, and naming conventions for SmartClone volumes.

Space requirements

SmartClone volumes inherit the size and data protection level of the source volume and snapshot. (When creating a SmartClone volume, you first create a snapshot of the source volume and create the SmartClone volumes from that snapshot, which is then called the “clone point.”) You can select the provisioning method when creating SmartClone volumes. See [“Provisioning storage” \(page 140\)](#) for a complete discussion of volume and snapshot characteristics and space planning.

- The space required for the volumes created using the SmartClone process is the same as for any other volumes on the SAN. SmartClone volumes can have schedules to snapshot a volume and remote snapshot a volume, just like other volumes, so the space requirements for SmartClone volumes should take into account the space needed for their local and remote snapshots.
- Number of SmartClone volumes—Plan the total number of SmartClone volumes you intend to create as part of your space requirements.

Note that you can create up to 25 SmartClone volumes as one operation in the HP P4000 Centralized Management Console, and then repeat the process to create the desired number of SmartClone volumes.

Use the CLI to create larger quantities of SmartClone volumes in a single operation.

- Thin or full provisioning—The type of provisioning you select affects the amount of space required on the SAN, just as it does for regular volumes.
- Data protection level—The data protection level of the source volume must be retained when creating SmartClone volumes, though you can change the data protection level after the SmartClone volumes are created. However, if you change the data protection level for any SmartClone volume, the data protection level for all replicated volumes automatically changes.

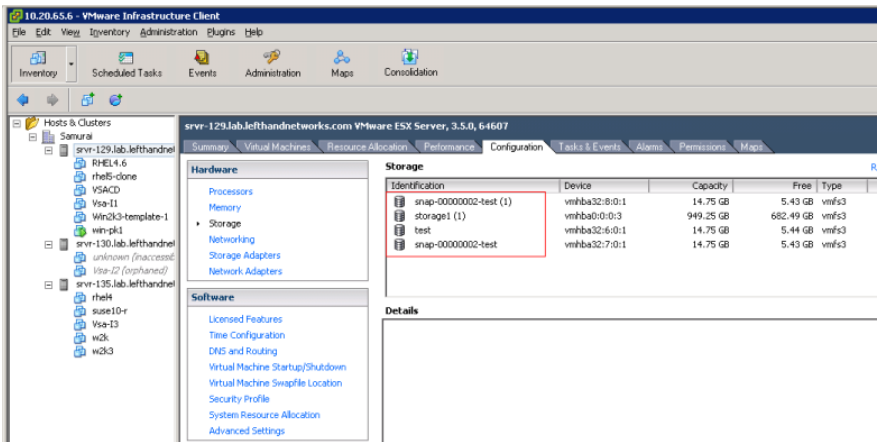
Naming convention for SmartClone volumes

A well-planned naming convention helps when you have many SmartClone volumes. Plan the naming ahead of time, since you cannot change volume or snapshot names after they have been created. You can design a custom naming convention when you create SmartClone volumes.

Naming and multiple identical disks in a server

Mounting multiple identical disks to servers typically requires that servers write new disk signatures to them. For example, VMware ESX servers require that resignaturing be enabled and will automatically name duplicate datastores. Most servers allow the duplicate disks to be renamed.

Figure 70 Duplicate names on duplicate datastores in ESX Server



Server access

Plan the servers you intend to assign to the SmartClone volumes. Configure the servers before creating the volumes, and you can then assign the servers when you create the volumes. See “Controlling server access to volumes” (page 196).

Defining SmartClone volume characteristics

When creating SmartClone volumes, you define the following characteristics.

Table 51 Characteristics for new SmartClone volumes

SmartClone volume characteristic	What it means
Quantity	The number of SmartClone volumes you want to create. You can create up to 25 as one operation in the CMC, and then repeat the process to create the desired number of SmartClone volumes. Use the CLI to create larger quantities of SmartClone volumes in a single operation.
SmartClone Name	The name of the SmartClone volume that is displayed in the CMC. A volume name is from 1 to 127 characters and is case sensitive. The name cannot be changed after the volume is created.
Provisioning	SmartClone volumes default to Thin Provisioning. You can select Full Provisioning when you create them. You can also edit the individual volumes after they are created and change the type of provisioning.
Server	Server assigned to the volume. While you can only assign one server when you create SmartClone volumes, you can go back and add additional clustered servers later. For

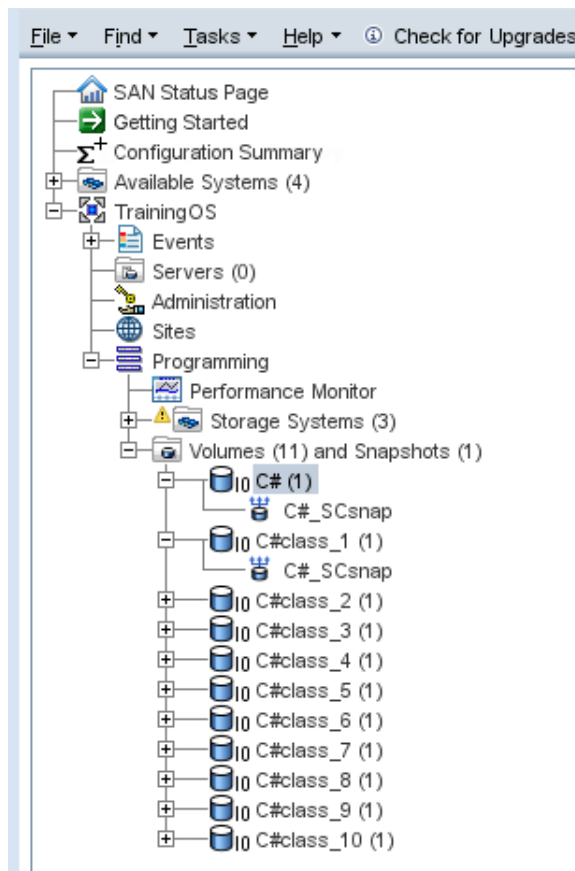
Table 51 Characteristics for new SmartClone volumes *(continued)*

SmartClone volume characteristic	What it means
	more information, see “Assigning server connections access to volumes” (page 202).
Permission	Type of access to the volume: Read, Read/Write, None

Naming SmartClone volumes

Because you may create dozens or even hundreds of SmartClone volumes, you need to plan the naming convention for them. For information about the default naming conventions built into the SAN/iQ software, see [“Setting naming conventions”](#) (page 15).

When you create a SmartClone volume, you can designate the base name for the volume. This base name is then used with numbers appended, incrementing to the total number of SmartClone volumes you create. For example, [Figure 71](#) (page 181) shows a SmartClone volume with the base name of “C#” and 10 clones. (The number in parentheses indicates how many snapshots are under that volume.)

Figure 71 Example of using a base name with 10 SmartClone volumes

After you designate a base name for the SmartClone volumes while you are creating them, you can then edit individual names of SmartClone volumes in the table list, before you finish creating them.

NOTE: Rename the SmartClone volume at the bottom of the list. Then the numbering sequence is not disrupted.

Figure 72 Rename SmartClone volume from base name

New SmartClone Volumes

Original Volume Setup

Management Group: TrainingOS

Volume Name: C#

Snapshot Name: C#_SCsnap New Snapshot...

SmartClone Volume Setup

Base Name: C#class **Provisioning:** Thin

Server: [No Server] **Permission:** Read/Write

Quantity (Max of 25): 10 Update Table

SmartClone Volume Name	Provisioning	Server Name	Permission
C#class_1	Thin	[No Server]	Read/Write
C#class_2	Thin	[No Server]	Read/Write
C#class_3	Thin	[No Server]	Read/Write
C#class_4	Thin	[No Server]	Read/Write
C#class_5	Thin	[No Server]	Read/Write
C#class_6	Thin	[No Server]	Read/Write
C#class_7	Thin	[No Server]	Read/Write
C#class_8	Thin	[No Server]	Read/Write
C#class_9	Thin	[No Server]	Read/Write
C#class_Beginner	Thin	[No Server]	Read/Write

OK **Cancel**

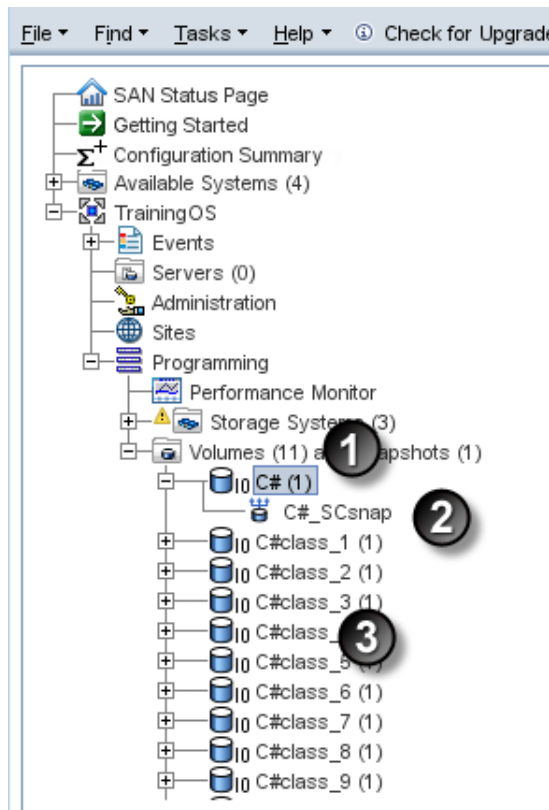
1. Rename SmartClone volume in list

Shared versus individual characteristics

Characteristics for SmartClone volumes are the same as for regular volumes. However, certain characteristics are shared among all the SmartClone volumes and snapshots created from a common clone point. If you want to change one of these shared characteristics for one SmartClone volume, that change will apply to all related volumes and snapshots, including the original volume and snapshot from which you created the SmartClone volumes. Simply use Edit Volume on the selected volume, and make the change to the volume. A message opens, stating that the change will apply to all of the related volumes, which are noted in the message.

For example, in [Figure 73 \(page 183\)](#), in the cluster Programming, there are 10 SmartClone volumes created from one source volume and its clone point. You want to move the first of the SmartClone volumes, to the cluster SysAdm2.

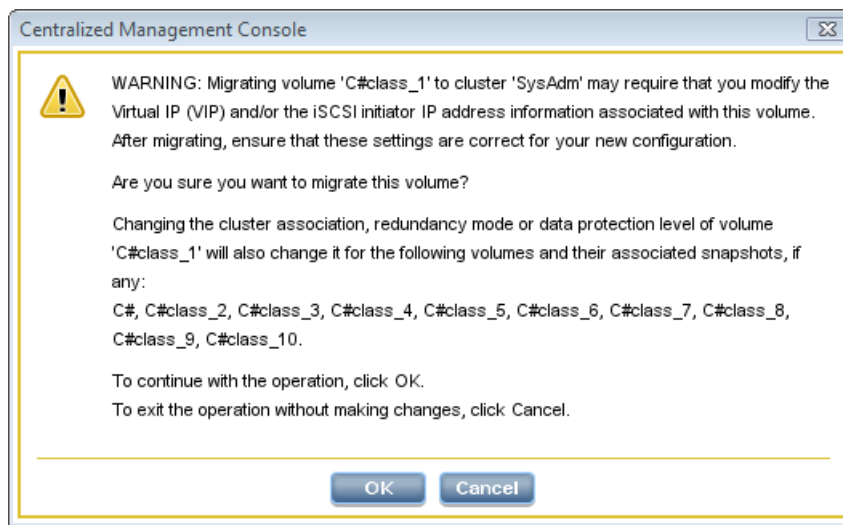
Figure 73 Programming cluster with SmartClone volumes, clone point, and the source volume



1. Source volume
2. Clone point
3. SmartClone volumes (5)

In this example, you edit the SmartClone volume, and on the Advanced tab you change the cluster to SysAdm. The confirmation message lists all the volumes and snapshots that will change clusters as a result of changing the edited volume.

Figure 74 Changing one SmartClone volume changes all related volumes and snapshots



When you click **OK** on the message, all the volumes and snapshots move to the cluster SysAdm.

Figure 75 SysAdm cluster now has the SmartClone volumes, clone point, and the source volume

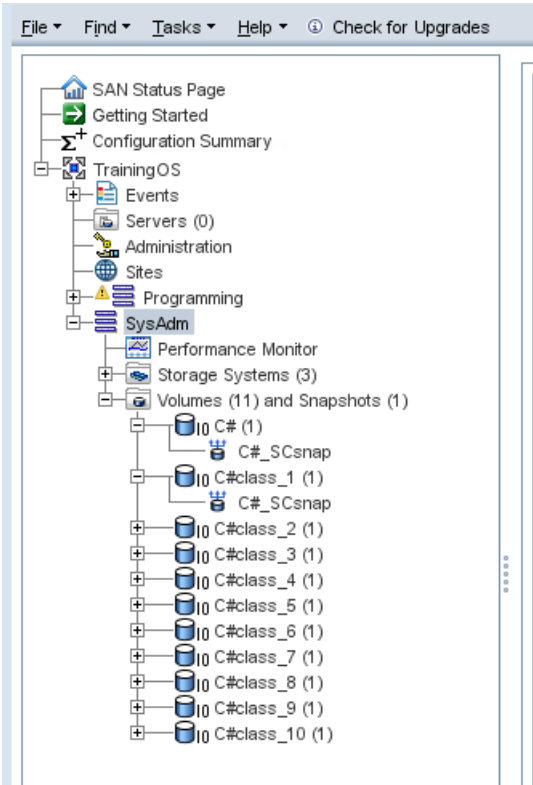


Table 52 (page 184) shows the shared and individual characteristics of SmartClone volumes. Note that if you change the cluster or the data protection level of one SmartClone volume, the cluster and data protection level of all the related volumes and snapshots will change.

Table 52 Characteristics of SmartClone volumes

Shared characteristics	Individual characteristics
Cluster	Name
Data protection level	Description
	Size
	Type (Primary or Remote)
	Provisioning (Thin or Full)
	Server

NOTE: Snapshot schedules and remote copy schedules are also individual to a single SmartClone volume.

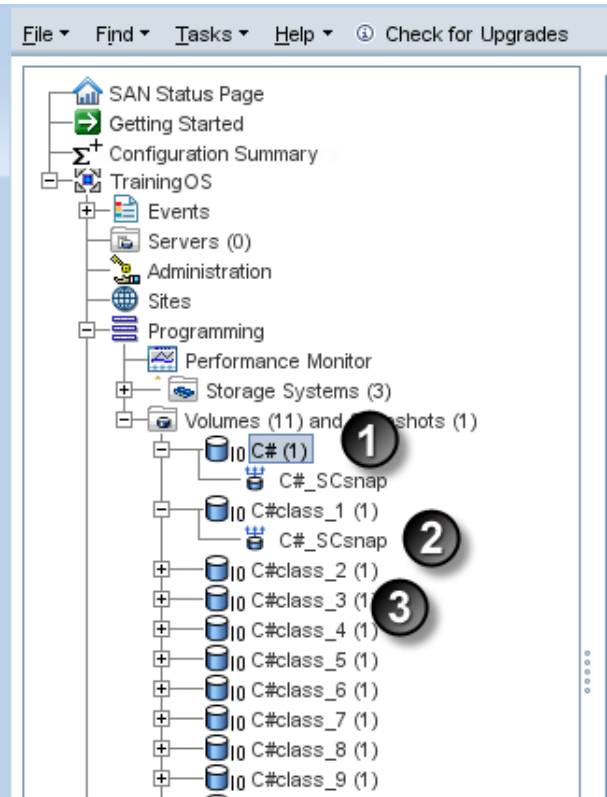
Clone point



The icon shown here represents the clone point in the navigation window. The clone point is the snapshot from which the SmartClone volumes are created. The clone point contains the snapshot data that is shared among the multiple volumes. Because the SmartClone volumes and their snapshots depend on the clone point, it cannot be deleted until it is no longer a clone point. A clone point ceases to be a clone point when only one SmartClone volume remains that was created from that

clone point. That is, you can delete all but one of the SmartClone volumes, and then you can delete the clone point.

Figure 76 Navigation window with clone point




- 1. Original volume
- 2. Clone point
- 3. SmartClone volume

In [Figure 76](#) (page 185), the original volume is “C#.”

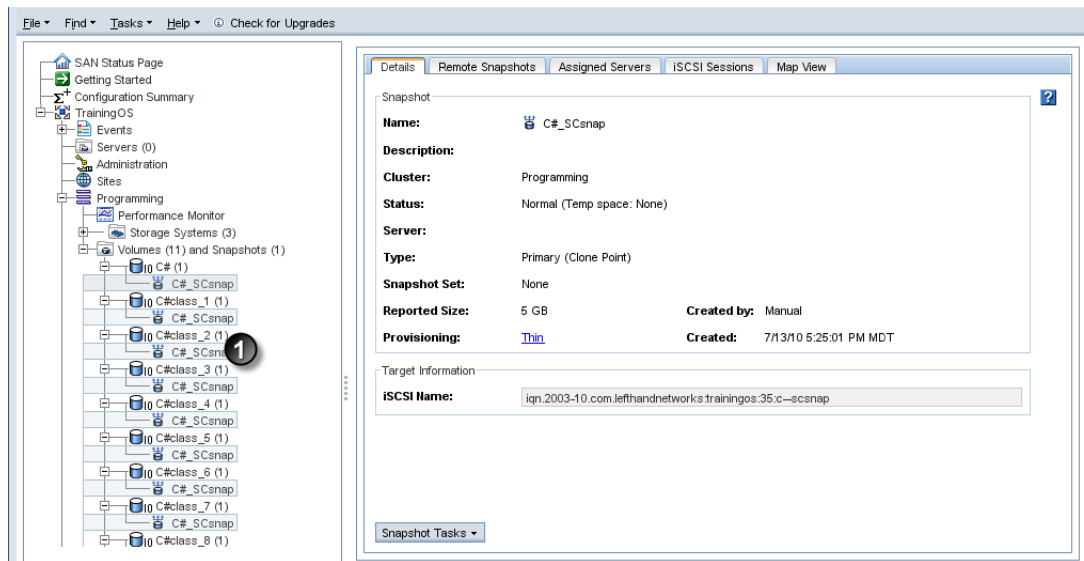
- Creating a SmartClone volume of C# first creates a snapshot, C#_SCsnap.
- After the snapshot is created, you create at least one SmartClone volume, C#class_1.

Table 53 How it works—clone point

First, a volume	C#
Next, a snapshot	C#_SCsnap
Next, SmartClone from the snapshot	C#class_1
Snapshot becomes a clone point	

Because the SmartClone volumes depend on the clone point from which they were created, the clone point appears underneath each SmartClone volume in the navigation window. While the clone point may appear many times, it only exists as a single snapshot in the SAN. Therefore, it only uses the space of that single snapshot. The display in the navigation window depicts this by the multiple highlights of the clone point underneath each SmartClone volume that was created from it.

Figure 77 Clone point appears under each SmartClone volume



1. Clone point appears multiple times. Note that it is exactly the same in each spot

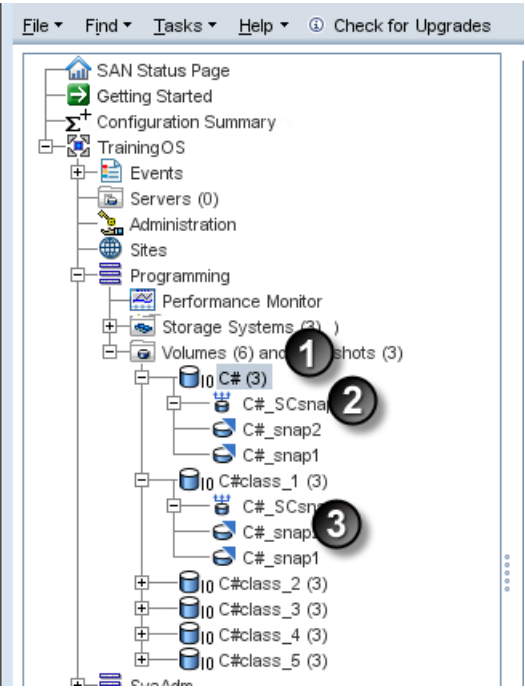
NOTE: Remember that a clone point only takes up space on the SAN once.

Shared snapshot



Shared snapshots occur when a clone point is created from a newer snapshot that has older snapshots below it in the tree. They are designated in the navigation window with the icon shown here.

Figure 78 Navigation window with shared snapshots





1. Original volume
2. Clone point
3. Shared snapshots

In [Figure 78 \(page 187\)](#), the original volume is C#. Three snapshots were created from C#:

- C#_snap1
- C#_snap2
- C#_SCsnap

Then a SmartClone volume was created from the latest snapshot, C#_SCsnap. That volume has a base name of C#_class. The older two snapshots, C#_snap1 and C#_snap2, become shared snapshots, because the SmartClone volume depends on the shared data in both those snapshots.

Table 54 How it works—shared snapshots

First, a volume	C#
Next, 3 snapshots	C#_snap1C#_snap2C#_SCsnap
Finally, SmartClone volumes from the latest snapshot	C#_class_x
latest snapshot becomes clone point	
Older two snapshots become shared between clone point and SmartClone volume.	

The shared snapshots also appear under all the volumes which share them. In [Figure 78 \(page 187\)](#), they are displayed under the original volume from which they were created, and under the single SmartClone volume that shares them. The selected shared snapshot is highlighted in the navigation window, under both the volumes with which it is shared. Shared snapshots can be deleted.

Creating SmartClone volumes

You create SmartClone volumes from existing volumes or snapshots. When you create a SmartClone volume from another volume, you first take a snapshot of the original volume. When you create a SmartClone volume from a snapshot, you do not take another snapshot.

To create a SmartClone volume

When you create SmartClone volumes, you either set the characteristics for the entire group or set them individually.

Figure 79 Setting characteristics for SmartClone volumes

New SmartClone Volumes

Original Volume Setup

Management Group: TrainingOS

Volume Name: C#

Snapshot Name: C#_SCsnap

SmartClone Volume Setup

Base Name: C#class Provisioning: Thin

Server: [No Server] Permission: Read/Write

Quantity (Max of 25): 1

Update Table

SmartClone Volume Name	Provisioning	Server Name	Permission
C#class_1	Thin	[No Server]	Read/Write

OK Cancel

1. Set characteristics for multiples here
2. Edit individual clones here

For details about the characteristics of SmartClone volumes, see [“Defining SmartClone volume characteristics”](#) (page 180).

1. Log in to the management group in which you want to create a SmartClone volume.
2. Select the volume or snapshot from which to create a SmartClone volume.
 - From the main menu you can select **Tasks→Volume→New SmartClone** or **Tasks→Snapshot→New SmartClone**.
Select the desired volume or snapshot from the list that opens.
 - In the navigation window, select the cluster and volume or snapshot from which to create a SmartClone volume.
3. Right-click on the volume or snapshot, and select **New SmartClone Volumes**.
4. If you are creating a SmartClone volume from a volume, click **New Snapshot** to first create a snapshot of the volume.

For more information, see [“Creating snapshots”](#) (page 163).

If you are creating a SmartClone volume from a snapshot, you do not create another snapshot.

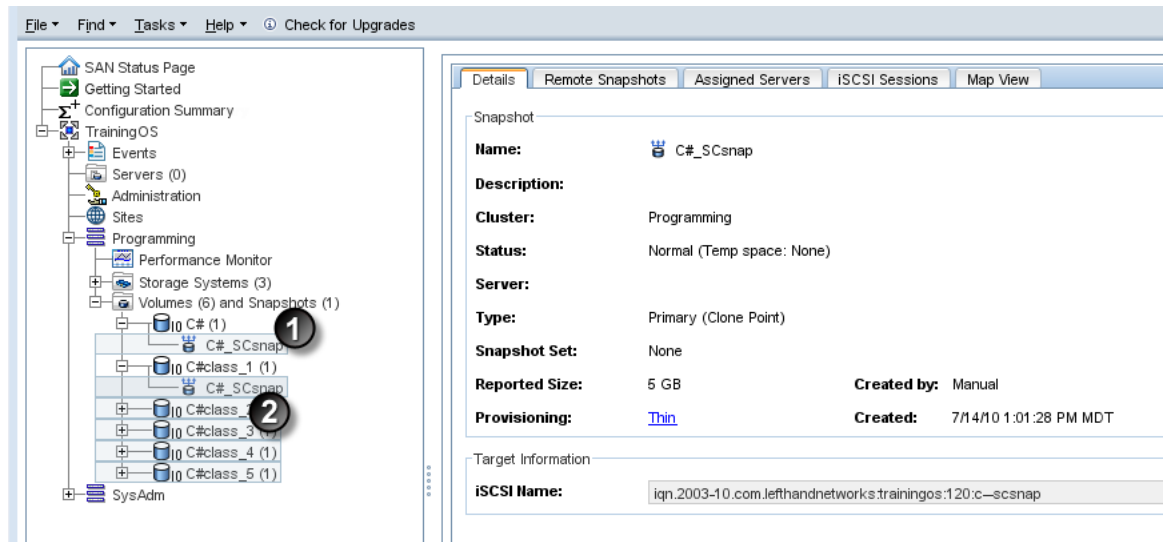
5. Next you select the following characteristics:
 - Base name for the SmartClone volumes
 - Type of provisioning
 - Server you want connected to the volumes, and
 - Appropriate permission.
6. In the Quantity field, select the number of SmartClone volumes you want to create.
7. Click **Update Table** to populate the table with the number of SmartClone volumes you selected.
8. If you want to modify any individual characteristic, do it in the list before you click **OK** to create the SmartClone volumes.

For example, you might want to change the assigned server of some of the SmartClone volumes. In the list you can change individual volumes' server assignments.

9. Click **OK** to create the volumes.

The new SmartClone volumes appear in the navigation window under the volume folder.

Figure 80 New SmartClone volumes in Navigation window



1. Clone point
2. New SmartClone volumes

Viewing SmartClone volumes

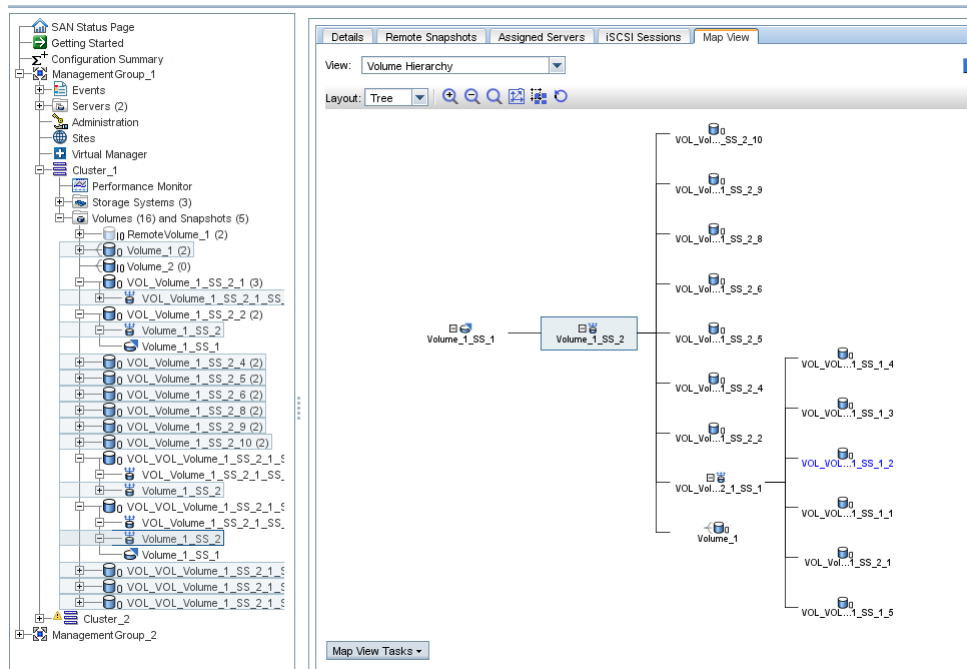
As you create multiple SmartClone volumes, you can view them and their related volumes and snapshots in both the navigation window and in the Map View tab, shown in [Figure 81 \(page 190\)](#).

Because a SmartClone volume is the same as any other volume, the icon is the standard volume icon. However, the clone point and the shared snapshot have unique icons, as illustrated in [Figure 76 \(page 185\)](#).

Map view

The Map View tab is useful for viewing the relationships between clone point snapshots, shared snapshots, and their related volumes. For example, when you want to make changes such as moving a volume to a different cluster, or deleting shared snapshots, the Map View tab allows you to easily identify how many snapshots and volumes are affected by such changes.

Figure 81 Viewing SmartClone volumes and snapshots as a tree in the Map View



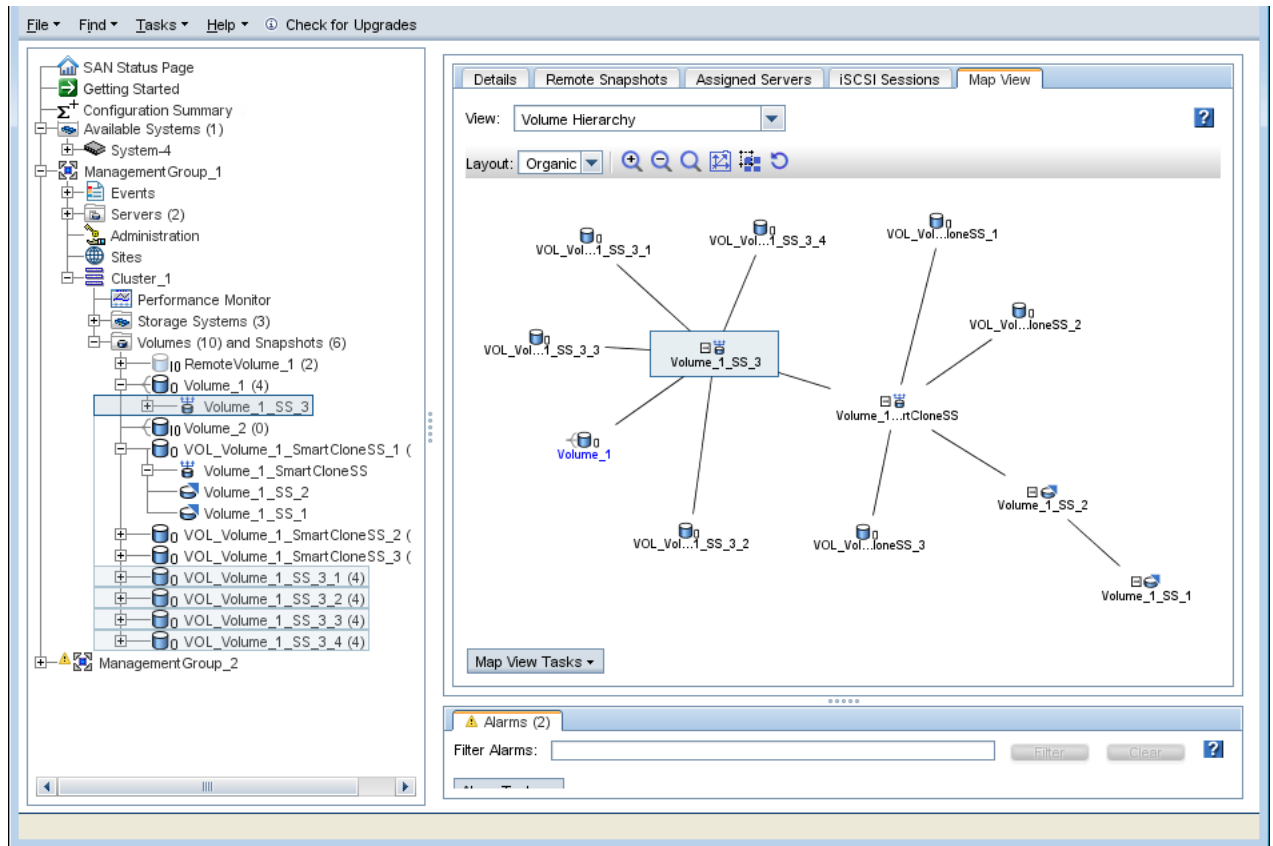
Using views

The default view is the tree layout, displayed in [Figure 81 \(page 190\)](#). The tree layout is the most effective view for smaller, more complex hierarchies with multiple clone points, such as clones of clones, or shared snapshots.

You may also display the Map view in the organic layout. The organic layout is more useful when you have a single clone point with many volumes, such as large numbers in a virtual desktop implementation. In such a case, the tree quickly becomes difficult to view, and it is much easier to distinguish the multiple volumes in the organic layout.

See [“Using the Map View ” \(page 14\)](#) for more information on using the map view display tools.

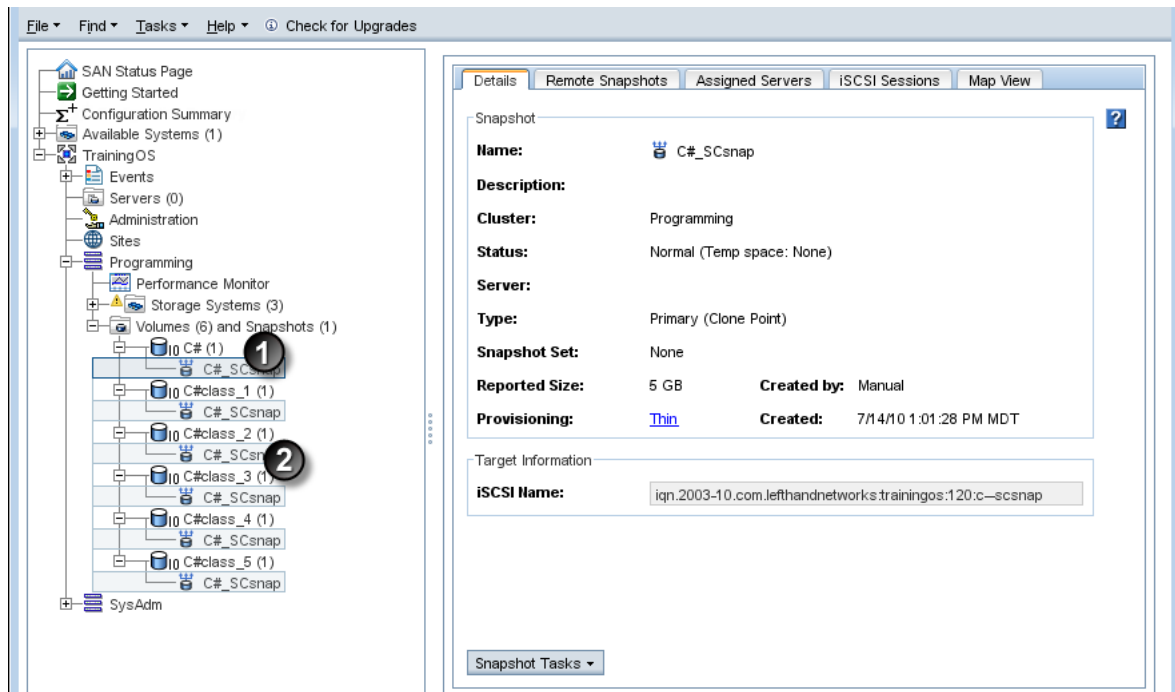
Figure 82 Viewing the organic layout of SmartClone volumes and related snapshots in the Map View



Viewing clone points, volumes, and snapshots

The navigation window view of SmartClone volumes, clone points, and snapshots includes highlighting that shows the relationship between related items. For example, in [Figure 83 \(page 192\)](#), the clone point is selected in the tree. The clone point supports the SmartClone volumes, so it is displayed under those volumes. The highlight shows the relationship of the clone point to the original volume plus the SmartClone volumes created from the original volume.

Figure 83 Highlighting all related clone points in navigation window



1. Selected clone point
2. Clone point repeated under SmartClone volumes

Editing SmartClone volumes

Use the Edit Volume window to change the characteristics of a SmartClone volume.

Table 55 Requirements for changing SmartClone volume characteristics

Item	Shared or Individual	Requirements for Changing
Description	Individual	May be up to 127 characters.
Size	Individual	Sets available space on cluster.
Servers	Individual	Existing server defined.
Cluster	Shared	<p>All related volumes and snapshots will move automatically to the target cluster. The target cluster must</p> <ul style="list-style-type: none"> Reside in the same management group. Have sufficient storage systems and unallocated space for the size and data protection level of the volume and all the other related volumes and snapshots being moved. <p>When moving volumes to a different cluster, those volumes temporarily exist on both clusters.</p>
Data protection Level	Shared	<p>All related volumes and snapshots must change to the same data protection level. The cluster must have sufficient storage systems and unallocated space to support the new data protection level for all related volumes.</p>

Table 55 Requirements for changing SmartClone volume characteristics *(continued)*

Item	Shared or Individual	Requirements for Changing
Type	Individual	Determines whether the volume is primary or remote.
Provisioning	Individual	Determines whether the volume is fully provisioned or thinly provisioned.

To edit the SmartClone volumes

1. In the navigation window, select the SmartClone volume for which you want to make changes.
2. Click **Volume Tasks**, and select **Edit Volume**.

See “Requirements for changing SmartClone volume characteristics” (page 192) for detailed information about making changes to the SmartClone volume characteristics.

3. Make the desired changes to the volume, and click **OK**.

If you change a SmartClone volume characteristic that will change other related volumes and snapshots, a warning opens that lists the volumes that will be affected by the change. If there are too many volumes to list, a subset will be listed with a note indicating how many additional volumes will be affected.

Deleting SmartClone volumes

Any volumes or snapshots that are part of a SmartClone network can be deleted just like any other volumes or snapshots. The only exception is the clone point, which cannot be deleted until it is no longer a clone point.

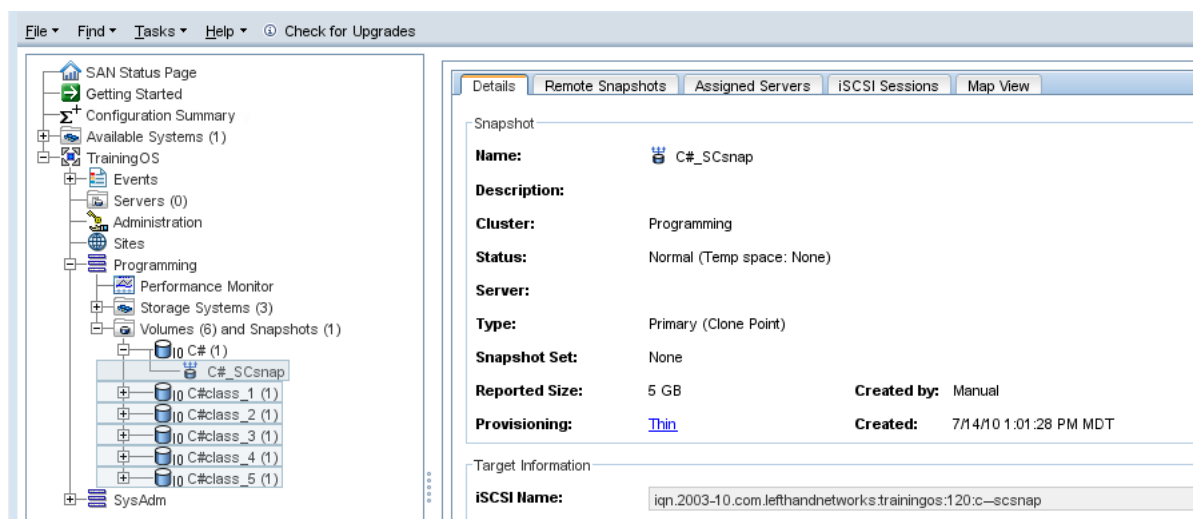
- ⚠ CAUTION:** Before deleting any volumes or snapshots, you must first stop any applications that are accessing the volumes and log off any iSCSI sessions that are connected to the volumes.

Deleting the clone point

You can delete a clone point if you delete all but one volume that depends on that clone point. After you delete all but one volume that depends on a clone point, the clone point returns to being a standard snapshot and can be managed just like any other snapshot.

For example, in [Figure 84 \(page 193\)](#), you must delete any four of the five C#class_x volumes before you can delete the clone point.

Figure 84 Viewing volumes that depend on a clone point

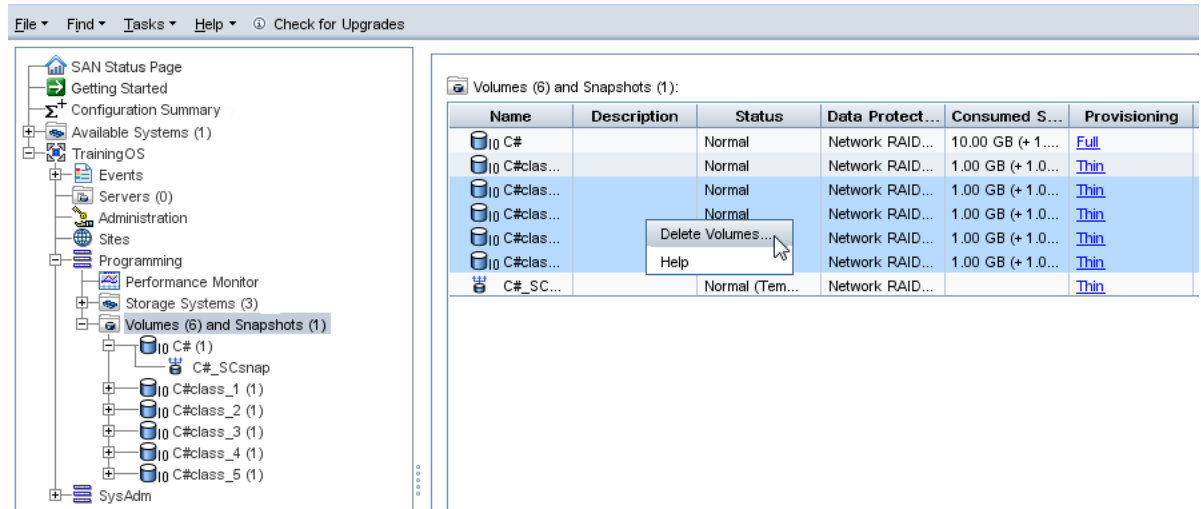


Deleting multiple SmartClone volumes

Delete multiple SmartClone volumes in a single operation from the Volume and Snapshots node of the cluster. First you must stop any application servers that are using the volumes, and log off any iSCSI sessions.

1. Select the **Volumes and Snapshots** node to display the list of SmartClone volumes in the cluster.

Figure 85 List of SmartClone volumes in cluster



2. Use **Shift+Click** to select the SmartClone volumes to delete.
3. Right-click, and select **Delete Volumes**.
A confirmation message opens.
4. When you are certain that you have stopped applications and logged off any iSCSI sessions, check the box to confirm the deletion, and click **Delete**.
It may take a few minutes to delete the volumes and snapshots from the SAN.

15 Working with scripting

Scripting in the SAN/iQ software through release 7.0 was accomplished by the `java.commandline.CommandLine` scripting tool.

In SAN/iQ software release 8.0, the `java.commandline.CommandLine` scripting tool was replaced by the HP P4000 Command-Line Interface. The CLI takes advantage of the new SAN/iQ API that provides comprehensive coverage of SAN/iQ software functionality to support scripting, integration and automation.

The `java.commandline.CommandLine` scripting tool will be supported after the 8.0 release to allow time for converting existing scripts that use `java.commandline.CommandLine` to the new CLI syntax.

Scripting documentation

- The *HP CLIQ - The SAN/iQ Command-Line Interface User Guide* is available from the HP website, and it is installed with the CLI.
- A SAN/iQ 8.0 Readme is available that describes the changes from `java.commandline.CommandLine` to the new CLI syntax.
- Sample scripts using the CLI are also available on the HP website.

16 Controlling server access to volumes

Application servers (servers), also called clients or hosts, access storage volumes on the SAN using the iSCSI protocol. You set up each server that needs to connect to volumes in a management group in the CMC. We refer to this setup as a “server connection.” Server connections can be single servers or server clusters. Use server clusters in the CMC to easily assign volumes to the clustered application servers that are accessing the SAN storage volumes. For example, create two servers and cluster them to assign two volumes to a Microsoft Exchange cluster or VMware ESX Server simultaneously.

You can set up servers to connect to volumes in three ways. All three ways use the virtual IP (VIP) for discovery and to log in to the volume from a server’s iSCSI initiator:

- iSCSI with VIP and load balancing—Use the load balancing option when you set up a server connection in the CMC to balance connections to the SAN.
- Microsoft DSM or HP P4000 DSM for MPIO (if using)—Automatically establishes multiple connections to the SAN.
- iSCSI with VIP only.

NOTE: Before setting up a server connection, make sure you are familiar with the iSCSI information in [“iSCSI and the HP P4000 SAN Solution”](#) (page 229).

Setting up server connections to volumes requires the general tasks outlined below.

Table 56 Overview of configuring server access to volumes

Task	For More Information
Ensure that an iSCSI initiator is installed on the server.	If you are using the HP P4000 DSM for MPIO, ensure that both the Microsoft MPIO and the DSM for MPIO are installed on the server. Refer to the <i>HP P4000 Application Integration Solution Pack User Guide</i> .
In the CMC, add the server connection to the management group and configure iSCSI access for that server.	See “Adding server connections to management groups” (page 197).
[Optional] Cluster multiple servers in the CMC to connect multiple volumes to clustered application servers.	See “Creating a server cluster” (page 200).
In the CMC, assign volumes to the server or server cluster.	See “Assigning server connections access to volumes” (page 202) “Assigning server connections access to volumes” (page 202).
In the iSCSI initiator on the server, log on to the volume.	See “Completing the iSCSI Initiator and disk setup” (page 204).
On the server, configure the volume using disk management tools.	See “Completing the iSCSI Initiator and disk setup” (page 204).

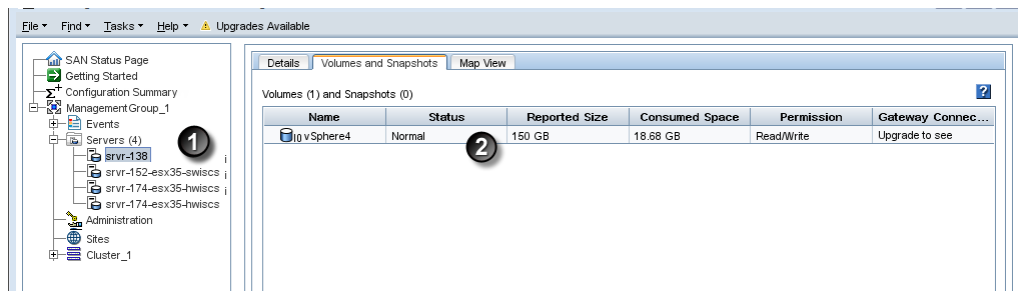
Change in server access control from version 7.0 and earlier

Before release 8.0, you controlled server access to volumes using authentication groups and volume lists. Starting with release 8.0, you control access with server and volume connections.

With release 8.0 and later, you add each server to a management group and assign server connections to volumes or snapshots. You can make the assignment from either the volume or the server.

The CMC displays the updated servers feature in management groups running SAN/iQ software version 7.x and earlier. However, these servers show the IQN number as their name.

Figure 86 Server assignments in the navigation window and the Volumes and Snapshots tab



1. Servers from a version 7.0 management group in the navigation window
2. Volumes and Snapshots tab shows the assigned volume that the server can access

Adding server connections to management groups

Add each server connection that needs access to a volume to the management group containing the volume. After you add a server connection to a management group, you can assign the server connection to one or more volumes or snapshots. For more information, see “Assigning server connections access to volumes” (page 202).

Prerequisites

- Each server must have an iSCSI initiator installed.
- The initiator node name, or `iqn` string, for the iSCSI initiator. See “iSCSI and CHAP terminology” (page 231).
- To use iSCSI load balancing, you must use a compliant iSCSI initiator. Verify the initiator compliance by going to the *HP P4000 SAN Solutions Compatibility Matrix* at: <http://www.hp.com/go/P4000compatibility>

CAUTION: Using a noncompliant iSCSI initiator with load balancing can compromise volume availability during iSCSI failover events.

Guide for servers

For detailed information about using iSCSI with the P4000 SAN Solution, including load balancing and CHAP authentication, see “iSCSI and the HP P4000 SAN Solution” (page 229).

Table 57 Characteristics for servers

Item	Description and requirements
Name	Name of the server that is displayed in the CMC. The server name is case sensitive and cannot be changed after the server is created.
Description	[Optional] Description of the server.
Controlling Server IP Address	Not used for the Application Aware Snapshot Manager on Windows. Required for the Application Aware Snapshot Manager on VMware. IP address of the Microsoft Windows server that is hosting the vCenter Server. Optional for using VSS on Windows servers.
Allow access via iSCSI	Allows the server to connect to the client using an iSCSI initiator.

Table 57 Characteristics for servers *(continued)*

Item	Description and requirements
Initiator Node Name	The name, or iqn string, of the iSCSI initiator. Open the iSCSI initiator and look for the string there. You can copy the string and paste it into the field.
Enable load balancing	Configures iSCSI load balancing for the server connections.
CHAP not required	Enables initiators to log on to a volume without authenticating their identity.
CHAP required	Requires initiators to authenticate before they can log on to a volume.

Adding a server connection

1. In the navigation window, log in to the management group.
2. Click **Management Group Tasks**, and select **New Server**.
3. Enter a name and description (optional) for the server connection.
4. If you are taking VMware application-managed snapshots, enter the Controlling Server IP Address.
5. Select **Allow access via iSCSI**.
6. If you want to use iSCSI load balancing and your initiator is compliant, select **Enable load balancing**.
7. In the Authentication section, select **CHAP not required**.
If later, you decide you want to use CHAP, you can edit the server connection (see [“Editing server connections”](#) (page 199)). For more information, see [“Authentication \(CHAP\)”](#) (page 230).
8. In the **Initiator Node Name** field, enter the iqn string.
9. Click **OK**.
10. [Optional] To use CHAP, edit the server connection you just configured and complete the fields necessary for the type of CHAP you intend to configure, as shown in [Table 58](#) (page 198).

Table 58 Entering CHAP information in a new server

For this CHAP Mode	Complete these fields
1-way CHAP	<ul style="list-style-type: none"> • CHAP name • Target secret—minimum of 12 characters
2-way CHAP	<ul style="list-style-type: none"> • CHAP name • Target secret—minimum of 12 characters • Initiator secret—minimum of 12 characters; must be alphanumeric

11. Click **OK**.
The server connection appears in the management group in the navigation window.
You can now assign this server connection to volumes, giving the server access to the volumes.
For more information, see [“Assigning server connections access to volumes”](#) (page 202).

Managing server connections

Manage your server connections:

- [“Editing server connections”](#) (page 199)
- [“Deleting server connections”](#) (page 199)

Editing server connections

You can edit the following fields for a server connection:

- Description
- Controlling Server IP Address
- Load balancing
- CHAP options

CAUTION: If you change the load balancing or CHAP options, you must log off and log back on to the target in the iSCSI initiator for the changes to take effect.

You can also delete a server connection from the management group. For more information, see [“Deleting server connections” \(page 199\)](#).

CAUTION: Editing the `iqn` or CHAP server settings may interrupt access to volumes. If necessary, or if the server is sensitive to disconnections, stop server access before editing a server.

1. In the navigation window, select the server connection you want to edit.
2. Click the **Details** tab.
3. Click **Server Tasks**, and select **Edit Server**.
4. Change the appropriate information.
5. Click **OK** when you are finished.
6. If you have changed the Enable Load Balancing option, a message opens, notifying you to log servers off and back on to the volumes.

This may entail stopping the applications, disconnecting them, reconnecting the applications to the volumes, and then restarting them.

Deleting server connections

Deleting a server connection stops access to volumes by servers using that server connection. Access to the same volume by other servers continues.

1. In the navigation window, select the server connection you want to delete.
2. Click the **Details** tab.
3. Click **Server Tasks**, and select **Delete Server**.
4. Click **OK** to delete the server.

Clustering server connections

You can cluster servers to easily assign multiple server connections to multiple volumes in a single operation. Cluster existing servers or you can create new servers to add to the server cluster from the New Server Cluster window. If you cluster servers that already have volumes assigned to them, all the servers in the cluster gain access to all the volumes, and inherit the access permissions of the original pair.

All clustered servers must have the same load balancing and volume access permission settings. You can verify and update these settings when you create the server cluster. For information about volume access permissions, see [“Server connection permission levels” \(page 203\)](#). For more information about iSCSI load balancing, see [“iSCSI load balancing” \(page 230\)](#).

Requirements for clustering servers

- Minimum of two servers for a cluster
- Same load balancing setting for all servers
- Same access level settings for all volumes assigned to the server cluster

Creating a server cluster

1. In the navigation window, select the Servers category.
2. Right-click and select **New Server Cluster**.
3. Enter a name and description (optional) for the server cluster.
4. Do one of the following:
 - Click **Add Server** and select the server from the list of available servers that opens.
 - Click **New Server** and follow the instructions found in [“Adding server connections to management groups”](#) (page 197).
5. Add the additional servers required for the server cluster, using either of the above methods.
6. [Optional] Next, click **Edit Server Cluster Settings** to verify that the server cluster settings are the same for all servers and volumes.

NOTE: The Server cluster Settings window opens automatically if inconsistencies are detected in the settings for the servers and volumes.

7. On the Server cluster Settings window, choose the proper settings for the server cluster. See [Figure 87](#) (page 200).
 - a. Select the appropriate radio button for the load balancing setting on each server.
 - b. Ensure that each volume listed has the same access permissions.

Figure 87 Verify the server cluster settings

Server cluster ServerCluster_1 settings

The load balancing and volume access permission settings must be the same for all servers in the server cluster. Select the desired settings below.

Load balancing

Load balancing must be the same for each server in the server cluster. Select the load balancing setting as appropriate.

☒ Enable load balancing on all servers in the server cluster.

☐ Do NOT enable load balancing on any server in the server cluster.

Volume access permissions

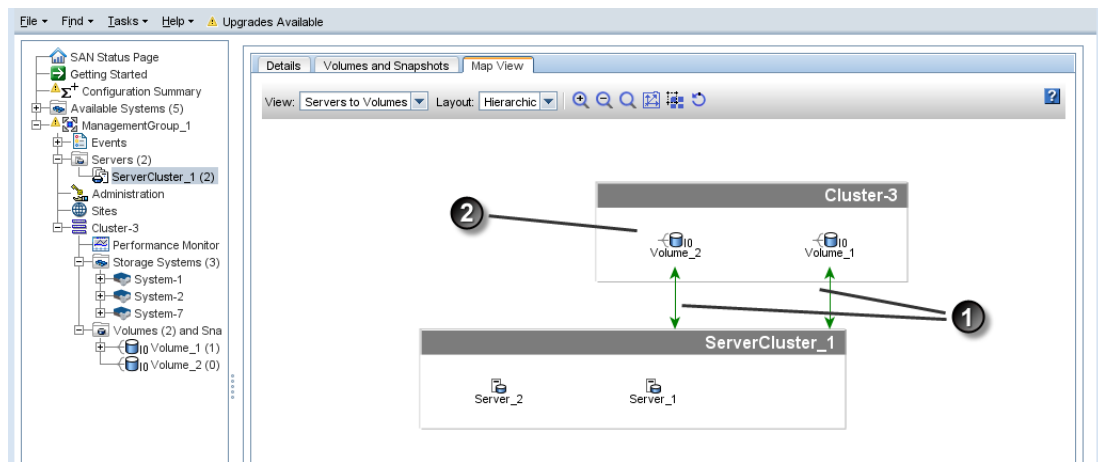
All servers in the server cluster must have the same access permissions to each volume. Select the access permission for all the servers in the server cluster for each of the volumes below.

Volume Name	Permission
Volume_2	Read/Write
Volume_1	Read/Write

OK Cancel

When the server cluster is created, you can view the results on the Servers Details tab and on Map View tab shown in [Figure 88](#) (page 201).

Figure 88 Completed server cluster and the assigned volumes



1. Green solid line indicates active connection. The two-way arrows indicate the volume permission levels are read-write. Black dotted line indicates an inactive iSCSI session.
2. Volumes have active iSCSI sessions, indicated by the line to the left of the volume icon

Server cluster map view

After you create a server cluster and connect volumes, use the Map View tab for viewing the relationships between systems, volumes and servers. For more information on using the map view tools, see [“Using the display tools”](#) (page 14).

Server cluster map views include the following:

- Servers to volumes
- Servers to systems

Working with a server cluster

Make any of the following changes to a server cluster:

- Add or remove servers.
- Edit a server to change iSCSI access or authentication, including settings for CHAP.
- Edit the server cluster settings for load balancing and volume access permissions. See [“Clustering server connections”](#) (page 199) for more information.

Editing a server cluster

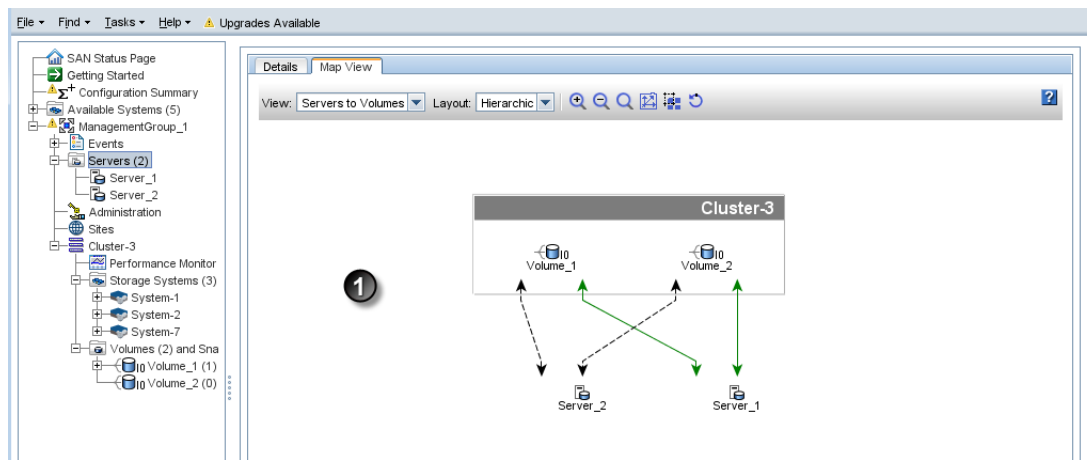
The Edit Server Cluster window includes sections for adding, editing, or removing servers, and managing load balancing and volume access permissions.

1. In the navigation window, select the Servers category, and select the server cluster to edit.
2. Right-click on the server cluster, and select **Edit Server Cluster**.
3. Make the desired changes to the server cluster in the Edit Server Cluster window.
4. Click **OK** when you have finished making the changes.

Deleting a server cluster

Deleting a server cluster removes the cluster associations among the servers. However, the associations remain between all of the volumes and servers that were in the cluster, shown in [Figure 89](#) (page 202). You must manually change the server and volume associations to the desired configuration after deleting the server cluster.

Figure 89 Servers and volumes retain connections after server cluster is deleted



1. Each volume remains connected to each server after the server cluster is deleted

To delete a server cluster and remove connections:

1. In the navigation window, select the Servers category and select the server cluster to delete.
2. Right-click on the server cluster and select **Delete Server Cluster**.
3. Select a server to change associations.
4. Right-click and select **Assign and Unassign Volumes and Snapshots**.
5. Select the appropriate volume and permission level for that server, and click **OK**.
6. Repeat these steps for the remaining servers, until the appropriate server and volume connections are in place.

Assigning server connections access to volumes

After you add a server connection to your management group, you can assign one or more volumes or snapshots to the server connection, giving the server access to those volumes or snapshots.

CAUTION: Without the use of shared storage access (host clustering or clustered file system) technology, allowing more than one iSCSI application server to connect to a volume at the same time, without cluster-aware applications and/or file systems in read/write mode, could result in data corruption.

You can make the assignments in two ways:

- “Assigning server connections from a volume” (page 203)
- “Assigning volumes from a server connection” (page 203)

Prerequisites

- The server connections you want to assign must already exist in the management group. See “Adding server connections to management groups” (page 197).
- The volumes or snapshots you want to assign must already exist in the management group.

When you assign the server connections and volumes or snapshots, you set the permissions that each server connection will have for each volume or snapshot. The available permissions are described in Table 59 (page 203).

Table 59 Server connection permission levels

Type of Access	Allows This
No access	Prevents the server from accessing the volume or snapshot.
Read access	Restricts the server to read-only access to the data on the volume or snapshot.
Read/write access	Allows the server read and write permissions to the volume.

NOTE: Microsoft Windows requires read/write access to volumes.

Assigning server connections from a volume

You can assign one or more server connections to any volume or snapshot. For the prerequisites, see [“Assigning server connections access to volumes” \(page 202\)](#).

1. In the navigation window, right-click the volume you want to assign server connections to.
2. Select **Assign and Unassign Servers**.
3. Select the **Assigned** check box for each server connection you want to assign to the volume or snapshot.
4. From the **Permission** list, select the permission each server connection should have to the volume or snapshot.
5. Click **OK**.

You can now log on to the volume from the server’s iSCSI initiator. See [“Completing the iSCSI Initiator and disk setup” \(page 204\)](#).

Assigning volumes from a server connection

You can assign one or more volumes or snapshots to any server connection. For the prerequisites, see [“Assigning server connections access to volumes” \(page 202\)](#).

1. In the navigation window, right-click the server connection you want to assign.
2. Select **Assign and Unassign Volumes and Snapshots**.
3. Select the **Assigned** check box for each volume or snapshot you want to assign to the server connection.
4. From the **Permission** list, select the permission the server should have.
5. Click **OK**.

You can now connect to the volume from the server’s iSCSI initiator. See [“Completing the iSCSI Initiator and disk setup” \(page 204\)](#).

Editing server connection and volume assignments

You can edit the assignment of volumes and server connections to:

- Unassign the volume or server connection
- Change the permissions



CAUTION: If you are going to unassign a server connection or restrict permissions, stop any applications from accessing the volume or snapshot, and log off the iSCSI session from the host before making the change.

Editing server connection assignments from a volume

You can edit the assignment of one or more server connections to any volume or snapshot.

1. In the navigation window, right-click the volume whose server connection assignments you want to edit.
2. Select **Assign and Unassign Servers**.
3. Change the settings as needed.
4. Click **OK**.

Editing server assignments from a server connection

You can edit the assignment of one or more volumes or snapshots to any server connection.

1. In the navigation window, right-click the server connection you want to edit.
2. Select **Assign and Unassign Volumes and Snapshots**.
3. Change the settings as needed.
4. Click **OK**.

Completing the iSCSI Initiator and disk setup

After you have assigned a server connection to one or more volumes, you must configure the appropriate iSCSI settings on the server. For information about iSCSI, see “[iSCSI and the HP P4000 SAN Solution](#)” (page 229).

For more information about setting up volumes and iSCSI, see the operating system-specific documents in the HP P4000 Manuals page of the HP Business Support Center website.

<http://www.hp.com/support/manuals>

In the Storage section, navigate to **Disk Storage Systems**→**P4000 SAN Solutions**→**HP LeftHand P4000 SAN Solutions** or **HP P4000 G2 SAN Solutions**, depending upon your product.

Persistent targets or favorite targets

After you configure the iSCSI initiator, you can log on to the volumes. When you log on, select the option to automatically restore connections. This sets up persistent targets that automatically reconnect after a reboot.

For persistent targets, you also need to set up dependencies to ensure that the applications on the server start only after the iSCSI service starts and the sessions are connected.

HP P4000 DSM for MPIO settings

If you are using HP P4000 DSM for MPIO and your server has two NICs, select the **Enable multi-path** option when logging on to the volume, and log on from each NIC.

For more information about DSM for MPIO, refer to the *HP P4000 Application Integration Solution Pack User Guide*.

Disk management

You must also format, configure, and label the volumes from the server using the operating system's disk management tools.

17 Monitoring performance

The Performance Monitor provides performance statistics for iSCSI and storage system I/Os to help you and HP support and engineering staff understand the load that the SAN is servicing.

The Performance Monitor presents real-time performance data in both tabular and graphical form as an integrated feature in the CMC. The CMC can also log the data for short periods of time (hours or days) to get a longer view of activity. The data will also be available via SNMP, so you can integrate with your current environment or archive the data for capacity planning. See [“Setting up SNMP”](#) (page 92).

As a real-time performance monitor, this feature helps you understand the current load on the SAN to provide additional data to support decisions on issues such as the following:

- Configuration options (Would network bonding help me?)
- Capacity expansion (Should I add more storage systems?)
- Data placement (Should this volume be on my SATA or SAS cluster?).

The performance data does not directly provide answers, but will let you analyze what is happening and provide support for these types of decisions.

These performance statistics are available on a cluster, volume, and storage system basis, letting you look at the workload on a specific volume and providing data like throughput, average I/O size, read/write mix, and number of outstanding I/Os. Having this data helps you better understand what performance you should expect in a given configuration. Storage system performance data will allow you to easily isolate, for example, a specific storage system with higher latencies than the other storage systems in the cluster.

Prerequisites

- You must have a cluster with one or more storage systems and one or more volumes connected via iSCSI sessions.
- All storage systems in the management group must have SAN/iQ software version 8.0 or later installed. The management group version on the Registration tab must show 8.0.
- A server must be accessing a volume to read data, write data, or both.

Introduction to using performance information

The Performance Monitor can monitor dozens of statistics related to each cluster.

The following sections offer some ideas about the statistics that are available to help you manage your SAN effectively. These sections cover just a few examples of common questions and issues, but they are not an exhaustive discussion of the possibilities the Performance Monitor offers.

For general concepts related to performance monitoring and analysis, see [“Performance monitoring and analysis concepts”](#) (page 215).

What can I learn about my SAN?

If you have questions such as these about your SAN, the Performance Monitor can help:

- What kind of load is the SAN under right now?
- How much more load can be added to an existing cluster?
- What is the impact of my nightly backups on the SAN?
- I think the SAN is idle, but the drive lights are blinking a lot. What is causing that?


Generally, the Performance Monitor can help you determine:

- Current SAN activities
- Workload characterization
- Fault isolation

Current SAN activities example

This example shows that the Denver cluster is handling an average of more than 747 IOPS with an average throughput of more than 6 million bytes per second and an average queue depth of 31.76.


Figure 90 Example showing overview of cluster activity



Workload characterization example

This example lets you analyze the workload generated by a server (ExchServer-1) including IOPS reads, writes, and total and the average IO size.


Figure 91 Example showing volume's type of workload



Fault isolation example

This example shows that the Denver-1 storage system (dotted line pegged at the top of the graph) has a much higher IO read latency than the Denver-3 storage system. Such a large difference may be due to a RAID rebuild on Denver-1. To improve the latency, you can lower the rebuild rate.

Figure 92 Example showing fault isolation



What can I learn about my volumes?

If you have questions such as these about your volumes, the Performance Monitor can help:

- Which volumes are accessed the most?
- What is the load being generated on a specific volume?


The Performance Monitor can let you see the following:

- Most active volumes
- Activity generated by a specific server

Most active volumes examples


This example shows two volumes (DB1 and Log1) and compares their total IOPS. You can see that Log1 averages about 2 times the IOPS of DB1. This might be helpful if you want to know which volume is busier.

Figure 93 Example showing IOPS of two volumes



This example shows two volumes (DB1 and Log1) and compares their total throughput. You can see that Log1 averages nearly 18 times the throughput of DB1. This might be helpful if you want to know which volume is busier.


Figure 94 Example showing throughput of two volumes



Activity generated by a specific server example

This example shows the total IOPS and throughput generated by the server (ExchServer-1) on two volumes.

Figure 95 Example showing activity generated by a specific server



Planning for SAN improvements

If you have questions such as these about planning for SAN improvements, the Performance Monitor can help:

- Would enabling NIC bonding on the storage systems improve performance?
- Is the load between two clusters balanced? If not, what should I do?
- I have budget to purchase two new storage systems.
 - Which volumes should I move to them to improve performance?
 - Which cluster should I add them to?


The Performance Monitor can let you see the following:

- Network utilization to determine if NIC bonding on the storage systems could improve performance
- Load comparison of two clusters
- Load comparison of two volumes

Network utilization to determine if NIC bonding could improve performance example

This example shows the network utilization of three storage systems. You can see that Denver-1 averages more than 79% utilization. You can increase the networking capacity available to that storage system by enabling NIC bonding on the storage system. You can also spread the load out across the storage systems using iSCSI load balancing.

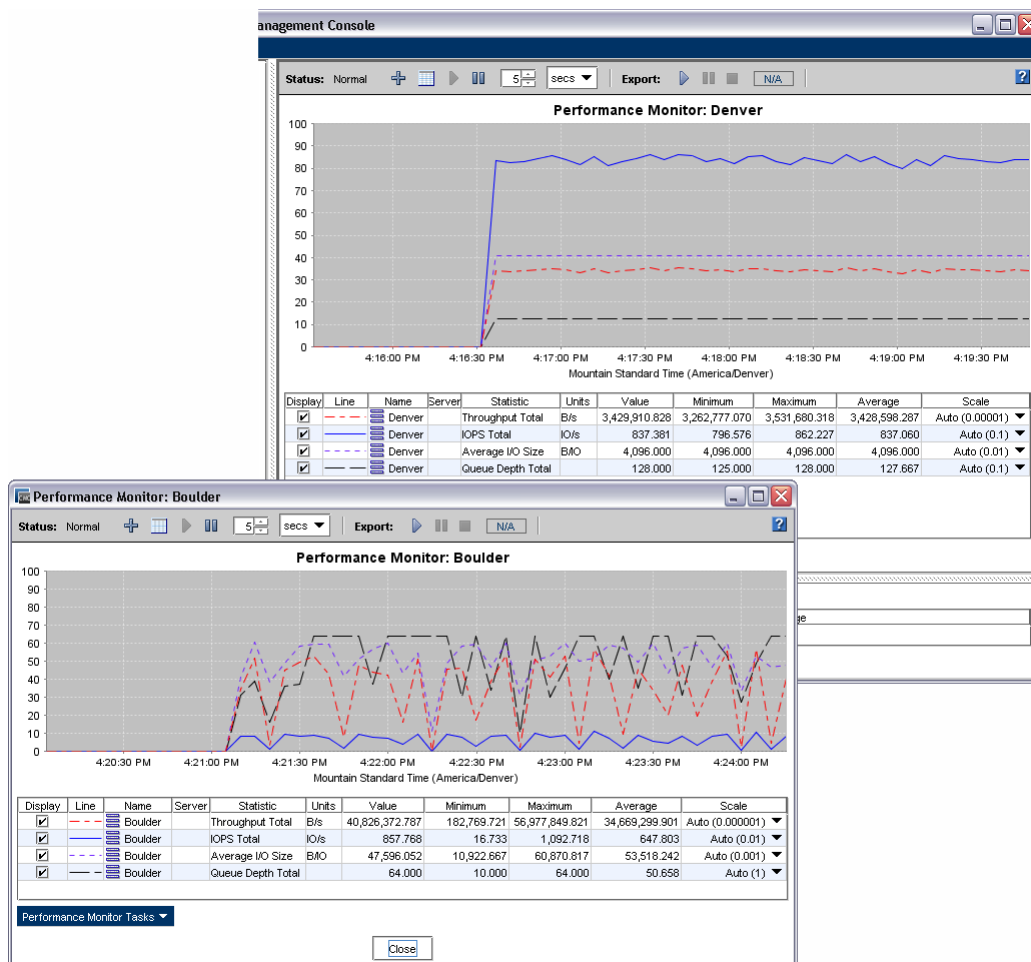
Figure 96 Example showing network utilization of three storage systems



Load comparison of two clusters example

This example illustrates the total IOPS, throughput, and queue depth of two different clusters (Denver and Boulder), letting you compare the usage of those clusters. You can also monitor one cluster in a separate window while doing other tasks in the CMC.


Figure 97 Example comparing two clusters



Load comparison of two volumes example

This example shows the total throughput for a cluster and the total throughput of each volume in that cluster. You can see that the Log1 volume generates most of the cluster's throughput.

Figure 98 Example comparing two volumes



Accessing and understanding the Performance Monitor window

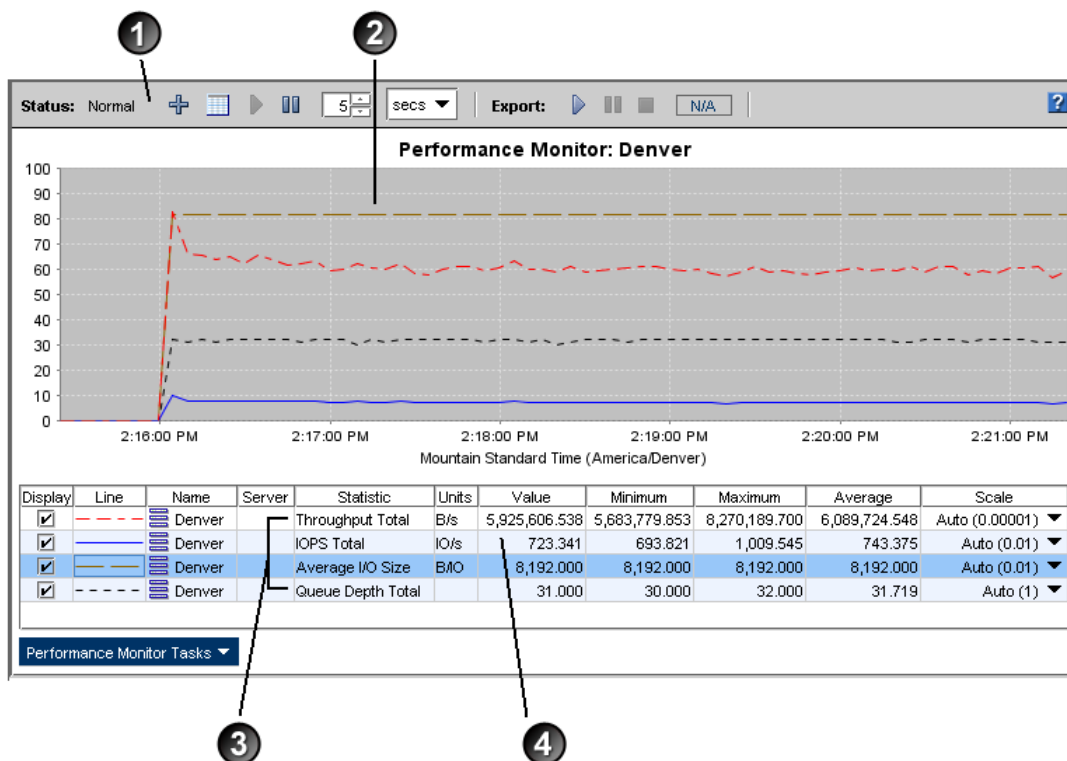
The Performance Monitor is available as a tree system below each cluster.

To display the Performance Monitor window:

1. In the navigation window, log in to the management group.
2. Select the Performance Monitor system for the cluster you want.

The Performance Monitor window opens. By default, it displays the cluster total IOPS, cluster total throughput, and cluster total queue depth.

Figure 99 Performance Monitor window and its parts



1. Toolbar
2. Graph
3. Default statistics
4. Statistics table

You can set up the Performance Monitor with the statistics you need. The system continues to monitor those statistics until you pause monitoring or change the statistics.

The system maintains any changes you make to the statistics graph or table only for your current CMC session. It reverts to the defaults the next time you log in to the CMC.

For more information about the performance monitor window, see the following:

- “Performance Monitor toolbar” (page 211)
- “Performance monitor graph” (page 211)
- “Performance monitor table” (page 212)

Performance Monitor toolbar

The toolbar lets you change some settings and export data.

Figure 100 Performance Monitor toolbar




Button or Status	Definition
1. Performance Monitor status	<ul style="list-style-type: none">• Normal—Performance monitoring for the cluster is OK.• Warning—The Performance Monitor is having difficulty monitoring one or more storage systems. Click the Warning text for more information.
2. Add Statistics	Opens the Add Statistics window.
3. Hide Graph/Show Graph	Toggles the graph display on or off.
4. Resume Monitoring	Restarts monitoring after pausing.
5. Pause Monitoring	Temporarily stops monitoring.
6. Sample Interval	Numeric value for the data update frequency.
7. Sample Interval Units	Unit of measure for the data update frequency, either minutes or seconds.
8. Export status	<ul style="list-style-type: none">• N/A—No export has been requested.• Sample interval and duration—If you have exported data, sample interval and duration display.• Paused—You paused an export.• Stopped—You stopped an export.• Warning—System could not export data. Click the Warning text for more information.• Error—System stopped the export because of a file IO error. Try the export again.
9. Start Export Log/Resume Export Log	Displays window to set up exporting of data to a comma separated values (CSV) file. Button changes to Resume Export Log when export is paused.
10. Pause Export Log	Temporarily stops exporting of data.
11. Stop Export Log	Stops the exporting of data.
12. Export Log Progress	Shows the progress of the current data export, based on the selected duration and elapsed time.

Performance monitor graph

The performance monitor graph shows a color-coded line for each displayed statistic.

Figure 101 Performance Monitor graph



The graph shows the last 100 data samples and updates the samples based on the sample interval setting.

The vertical axis uses a scale of 0 to 100. Graph data is automatically adjusted to fit the scale. For example, if a statistic value was larger than 100, say 4,000.0, the system would scale it down to 40.0 using a scaling factor of 0.01. If the statistic value is smaller than 10.0, for example 7.5, the system would scale it up to 75 using a scaling factor of 10. The Scale column of the statistics table shows the current scaling factor. If needed, you can change the scaling factor. For more information, see [“Changing the scaling factor”](#) (page 220).

The horizontal axis shows either the local time or Greenwich Mean Time. The default setting is the local time of the computer that is running the CMC. You can change this default to GMT. See [“Changing the sample interval and time zone”](#) (page 216). (This time zone setting is not related to the management group time zone.)

For information about controlling the look of the graph, see [“Changing the graph”](#) (page 219).

Performance monitor table

The performance monitor table displays a row for each selected statistic.

Figure 102 Performance Monitor table

Display	Line	Name	Server	Statistic	Units	Value	Minimum	Maximum	Average	Scale
<input checked="" type="checkbox"/>		Stores		Throughput Total	B/s	72,187,172.208	72,187,172.208	81,685,942.857	76,622,128.656	Auto (0.000001)
<input checked="" type="checkbox"/>		Stores		IOPS Total	IO/s	1,101.489	1,101.489	1,246.429	1,169.161	Auto (0.01)
<input checked="" type="checkbox"/>		Stores		Queue Depth Total		8.000	0.000	8.000	1.619	Auto (10)

The table shows information about the statistics selected for monitoring. The table values update based on the sample interval setting.

To view the statistic definition, hold your mouse pointer over a table row.

[Table 60](#) (page 212) defines each column of the Performance Monitor table.

Table 60 Performance Monitor table columns

Column	Definition
Display	Check box where you toggle the display of the graph line on or off.
Line	Shows the current color and style for the statistic's line on the graph.
Name	Name of the cluster, storage system, or volume being monitored.
Server	For volumes and snapshots, the server that has access.
Statistic	The statistic you selected for monitoring.

Table 60 Performance Monitor table columns *(continued)*

Column	Definition
Units	Unit of measure for the statistic.
Value	Current sample value for the statistic.
Minimum	Lowest recorded sample value of the last 100 samples.
Maximum	Highest recorded sample value of the last 100 samples.
Average	Average of the last 100 recorded sample values.
Scale	Scaling factor used to fit the data on the graph's 0 to 100 scale. Only the line on the graph is scaled; the values in the table are <i>not</i> scaled. The values in the log file, if you export the file, are also <i>not</i> scaled.

For information about adding statistics, see [“Adding statistics”](#) (page 216).

Understanding the performance statistics


You can select the performance statistics that you want to monitor.

For clusters, volumes, and snapshots, the statistics being reported are based on client IO. This is the iSCSI traffic and does not include other traffic such as replication, remote snapshots, and management functions.

For storage systems and devices, the statistics report the total traffic, including iSCSI traffic along with replication, remote snapshots, and management functions.

The difference between what the cluster, volumes, and snapshots are reporting and what the storage systems and devices are reporting is the overhead (replication, remote snapshots, and management functions).

Figure 103 Performance statistics and where they are measured



The following statistics are available:

Table 61 Performance Monitor statistics

Statistic	Definition	Cluster	Volume or Snapshot	NSM
IOPS Reads	Average read requests per second for the sample interval.	X	X	X
IOPS Writes	Average write requests per second for the sample interval.	X	X	X
IOPS Total	Average read+write requests per second for the sample interval.	X	X	X
Throughput Reads	Average read bytes per second for the sample interval.	X	X	X
Throughput Writes	Average write bytes per second for the sample interval.	X	X	X
Throughput Total	Average read and write bytes per second for the sample interval.	X	X	X
Average Read Size	Average read transfer size for the sample interval.	X	X	X
Average Write Size	Average write transfer size for the sample interval.	X	X	X
Average I/O Size	Average read and write transfer size for the sample interval.	X	X	X
Queue Depth Reads	Number of outstanding read requests.	X	X	-
Queue Depth Writes	Number of outstanding write requests.	X	X	-
Queue Depth Total	Number of outstanding read and write requests.	X	X	X
IO Latency Reads	Average time, in milliseconds, to service read requests.	X	X	X
IO Latency Writes	Average time, in milliseconds, to service write requests.	X	X	X
IO Latency Total	Average time, in milliseconds, to service read and write requests.	X	X	X
Cache Hits Reads	Percent of reads served from cache for the sample interval.	X	X	-
CPU Utilization	Percent of processor used on this storage	-	-	X

Table 61 Performance Monitor statistics *(continued)*

Statistic	Definition	Cluster	Volume or Snapshot	NSM
	system for the sample interval.			
Memory Utilization	Percent of total memory used on this storage system for the sample interval.	-	-	X
Network Utilization	Percent of bidirectional network capacity used on this network interface on this storage system for the sample interval.	-	-	X
Network Bytes Read	Bytes read from the network for the sample interval.	-	-	X
Network Bytes Write	Bytes written to the network for the sample interval.	-	-	X
Network Bytes Total	Bytes read and written over the network for the sample interval.	-	-	X
Storage Server Total Latency	Average time, in milliseconds, for the RAID controller to service read and write requests.	-	-	X

Monitoring and comparing multiple clusters

You can open the Performance Monitor for a cluster in a separate window. This lets you monitor and compare multiple clusters at the same time. You can open one window per cluster and rearrange the windows to suit your needs.

1. From the Performance Monitor window, right-click anywhere, and select **Open in Window**.
The Performance Monitor window opens as a separate window.
Use the **Performance Monitor Tasks** menu to change the window settings.
2. When you no longer need the separate window, click **Close**.

Performance monitoring and analysis concepts

The following general concepts are related to performance monitoring and analysis.

Workloads

A workload defines a specific characterization of disk activity. These characteristics are access type, access size, access pattern, and queue depth. Application and system workloads can be analyzed, then described using these characteristics. Given these workload characterizations, test tools like iometer can be used to simulate these workloads.

Access type

Disk accesses are either read or write operations. In the absence of disk or controller caching, reads and writes operate at the same speed.

Access size

The size of a read or write operation. As this size increases, throughput usually increases because a disk access consists of a seek and a data transfer. With more data to transfer, the relative cost of the seek decreases. Some applications allow tuning the size of read and write buffers, but there are practical limits to this.

Access pattern

Disk accesses can be sequential or random. In general, sequential accesses are faster than random accesses, because every random access usually requires a disk seek.

Queue depth

Queue depth is a measure of concurrency. If queue depth is one ($q=1$), it is called serial. In serial accesses, disk operations are issued one after another with only one outstanding request at any given time. In general, throughput increases with queue depth. Usually, only database applications allow the tuning of queue depth.

Changing the sample interval and time zone

You can set the sample interval to any value between 5 seconds and 60 minutes, in increments of either seconds or minutes.

The time zone comes from the local computer where you are running the CMC.

You can change the sample interval in the following ways:

- Using the toolbar
- In the Edit Monitoring Interval window, where you can also change the time zone

To change the sample interval from the toolbar:

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. In the toolbar, change the Sample Interval value.
4. In the toolbar, select the Sample Interval Units you want.

The Performance Monitor starts using the new interval immediately.

To change the sample interval and time zone:

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. Click **Performance Monitoring Tasks**, and select **Edit Monitoring Interval**.
4. In the Sample Every fields, enter the interval, and select the units you want.
5. Select **Local** or **Greenwich Mean Time**.
6. Click **OK**.

The Performance Monitor starts using the new interval and time zone immediately.

Adding statistics

You can change the monitored statistics for the Performance Monitor as needed. To limit the performance impact on the cluster, you can add up to 50 statistics.

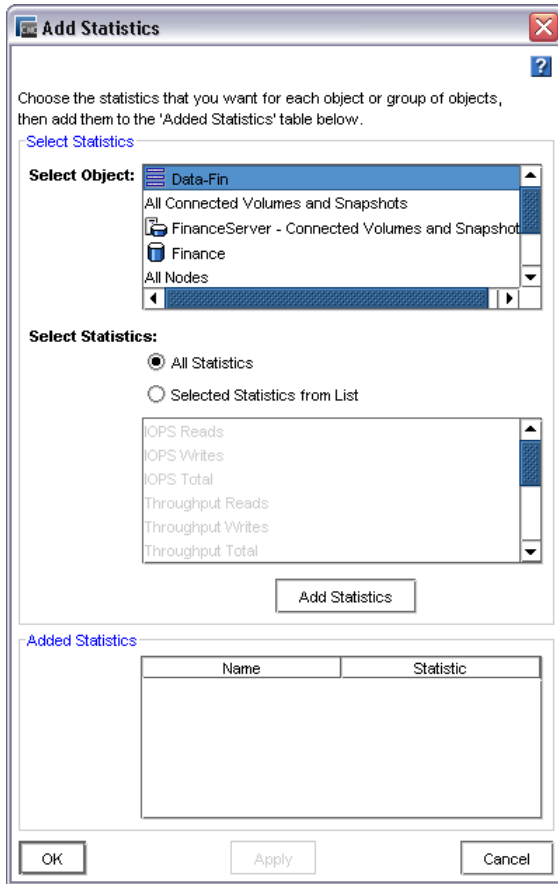
The system maintains any changes you make to the statistics only for your current CMC session. It reverts to the defaults the next time you log in to the CMC.

For definitions of the available statistics, see [“Understanding the performance statistics”](#) (page 213).

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.

3. Click .

Figure 104 Add Statistics window



4. From the Select Object list, select the cluster, volumes, and storage systems you want to monitor. Use the **CTRL** key to select multiple objects from the list.
5. From the Select Statistics options, select the option you want.
 - **All Statistics**—Adds all available statistics for each selected object.
 - **Selected Statistics from List**—Lets you select the statistics you want from the list below. The list is populated with the statistics that relate to the selected objects.
 Use the **CTRL** key to select multiple statistics from the list.
6. If you selected the Selected Statistics from List option, select the statistics you want to monitor.
7. Click **Add Statistics**.
If you selected a statistic that is already being monitored, a message appears letting you know that the statistics will not be added again.
8. Click **OK**.

Viewing statistic details

In addition to what you see in a table row, you can see all of the details for a specific statistic in the table, including the statistic definition.

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.

3. Right-click a row in the table and select **View Statistic Details**.

The Statistic Details window opens, with all of the information for the selected statistic that is in the table, plus the statistic definition.

4. Click **Close**.

Removing and clearing statistics

You can remove or clear statistics in any of the following ways:

- Remove one or more statistics from the table and graph
- Clear the sample data, but retain the statistics in the table
- Clear the graph display, but retain the statistics in the table
- Reset to the default statistics

Removing a statistic

You can remove one or more statistics from the table and graph.

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. Right-click a row in the table, and select **Remove Statistics**.
Use the **CTRL** key to select multiple statistics from the table.
4. Click **OK** to confirm.

Clearing the sample data

You can clear all the sample data, which sets all table values to zero and removes all lines from the graph. This leaves all of the statistics in the table and selected for display. The graph and table data repopulate with the latest values after the next sample interval elapses.

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. Right-click anywhere in the Performance Monitor window, and select **Clear Samples**.

Clearing the display

You can clear the display, which removes all lines from the graph and deselects the Display option for each statistic in the table. This leaves all of the statistics in the table, along with their data, which continue to update.

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. Right-click anywhere in the Performance Monitor window, and select **Clear Display**.



Resetting defaults

You can reset the statistics to the defaults, which removes all lines from the graph and sets the three default statistics (cluster total IOPS, cluster total throughput, and cluster total queue depth) to zero in the table. The default statistics are set to display and their data update when the next sample interval elapses.

1. In the navigation window, log in to the management group.
2. Select the **Performance Monitor** node for the cluster you want.
3. Right-click anywhere in the Performance Monitor window, and select **Reset to Defaults**.

Pausing and restarting monitoring

If you are currently monitoring one or more statistics, you can pause the monitoring and restart it. For example, you may want to pause monitoring during planned maintenance windows or production downtime.

1. From the **Performance Monitor** window, click  to pause the monitoring.
All data remain as they were when you paused.
2. To restart the monitoring, click .
Data updates when the next sample interval elapses. The graph will have a gap in the time.



Changing the graph

You can change the graph and its lines in the following ways:

- “Hiding and showing the graph” (page 219)
- “Displaying or hiding a line” (page 219)
- “Changing the color or style of a line” (page 219)
- “Highlighting a line” (page 219)
- “Changing the scaling factor” (page 220)

Hiding and showing the graph

By default, the performance monitor graph appears in the Performance Monitor window. If you want more space to display the performance monitor table, you can hide the graph.

1. From the Performance Monitor window, click  to hide the graph.
2. To redisplay the graph, click  to show the graph.

Displaying or hiding a line

When you add statistics to monitor, by default, they are set to display in the graph. You can control which statistics display in the graph, as needed.

1. From the Performance Monitor window, deselect the **Display** check box for the statistic in the table.
2. To redisplay the line, select the **Display** check box for the statistic.
If you want to display all of the statistics from the table, right-click anywhere in the Performance Monitor window, and select **Display All**.

Changing the color or style of a line

You can change the color and style of any line on the graph.

1. From the Performance Monitor window, select one or more statistics in the table that you want to change.
2. Right-click, and select **Edit Line**.
3. Select the color and line style options you want.
4. To see the changes and leave the window open, click **Apply**.
5. When you finish the changes, click **OK**.

Highlighting a line

You can highlight one or more lines on the graph to make them easier to distinguish.

1. From the Performance Monitor window, right-click the statistic in the table that you want to highlight, and select **Highlight**.
The line turns white.
2. To remove the highlight, right-click the statistic, and select **Remove Highlight**.

Changing the scaling factor

The vertical axis uses a scale of 0 to 100. Graph data is automatically adjusted to fit the scale. For example, if a statistic value was larger than 100, say 4,000.0, the system would scale it down to 40.0 using a scaling factor of 0.01. If the statistic value is smaller than 10.0, for example 7.5, the system would scale it up to 75 using a scaling factor of 10. The Scale column of the statistics table shows the current scaling factor. If needed, you can change the scaling factor. For example, if you are looking at similar items, you might change the scaling factor to change the emphasis on one item.

- From the statistics table on the Performance Monitor window, select the scaling factor you want from Scale drop-down list for the statistic you want to change.

The line moves up or down the graph based on the new scaling factor.

If the line is at the very top or bottom of the graph, the scaling factor is too large or too small to fit on the graph. It is possible for more than one line to be “pegged” to the top or bottom of the graph in this way, resulting in one or more lines being hidden behind another line. Set the Scale back to Auto to display the line.


Exporting data

You can export performance statistics to a CSV file or save the current graph to an image file.

Exporting statistics to a CSV file

You can export performance statistics to a CSV file. You select which statistics to export. They can be different from the statistics you are currently monitoring.

You also select the sample interval and the duration of the sampled data for export. Typical durations are from 10 minutes to 24 hours. The maximum duration is 999 hours, which is about 41 days.

1. From the Performance Monitor window, click  to start the export.
2. In the Log File field, enter the name of the file.

By default, the system saves the file to the My Documents folder (Windows) or your home directory (Linux) with a file name that starts with Performance and includes the cluster name, along with the date and time.

To select a different location, click **Browse**.

3. Set the Sample Every fields to the value and units you want for the sample interval.
4. Set the For Duration Of fields to the value and units you want for the monitoring period.
5. Click **Add Statistics**.
6. In the Add Statistics window, from the Select Object list, select the cluster, volumes, and storage systems you want to monitor.
Use the **CTRL** key to select multiple objects from the list.
7. From the Select Statistics options, select the option you want.
 - **All Statistics**—Adds all available statistics for each selected object.
 - **Selected Statistics from List**—Lets you select the statistics you want from the list below. The list is populated with the statistics that relate to the selected objects.
Use the **CTRL** key to select multiple statistics from the list.
8. If you selected the Selected Statistics from List option, select the statistics you want to monitor.

9. Click **Add Statistics**.



The statistics you selected are listed in the Added Statistics list.

10. Click **OK**.

The File Size field displays an estimated file size, based on the sample interval, duration, and selected statistics.

11. When the export information is set the way you want it, click **OK** to start the export.

The export progress appears in the Performance Monitor window, based on the duration and elapsed time.

To pause the export, click , then click  to resume the export.

To stop the export, click . Data already exported is saved in the CSV file.

Saving the graph to an image file

You can save the graph and the currently visible portion of the statistics table to an image file. This may be helpful if you are working with technical support or internal personnel to troubleshoot an issue.

1. From the Performance Monitor window, make sure the graph and table display the data you want.
2. Right-click anywhere in the Performance Monitor window, and select **Save Image**.
3. Navigate to where you want to save the file.
4. Change the file name, if needed.

The file name defaults to include the name of the object being monitored and the date and time.

5. Change the file type, if needed.
6. Click **Save**.

18 Registering advanced features

Advanced features expand the capabilities of the SAN/iQ software and are enabled by licensing the storage systems through the HP License Key Delivery Service website, using the license entitlement certificate that is packaged with each storage system. However, you can use the advanced features immediately by agreeing to enter an evaluation period when you begin using the SAN/iQ software for clustered storage. Throughout the evaluation period you receive reminders to register and obtain a license for the advanced features. Create the SAN and begin using all the available features before obtaining the licenses, and then apply the license keys when you obtain them.

The advanced features that require licensing are listed below.

- Multi-System Virtualization and Clustering—clustered storage systems that create a single pool of storage
- Managed Snapshots—recurring scheduled snapshots of volumes
- Remote Copy—scheduled or manual asynchronous replication of data to remote sites
- Multi-Site SAN—automatic synchronous data mirroring between sites

Evaluation period for using advanced features

When you use any feature that requires a license key, a message opens, asking you to verify that you want to enter the evaluation period. During this evaluation period you may configure, test, and modify any feature. At the end of the evaluation period, if you do not obtain a license key, all volumes and snapshots related to the feature become unavailable to any clients. The data is safe, and you can continue to manage the volumes and snapshots in the CMC. You can restore the entire configuration to availability by obtaining the license keys and applying them to the storage systems in the management group that contains the configured advanced features.

Starting the evaluation period

You start the evaluation period for an advanced feature when you configure that feature in the CMC.

Table 62 Descriptions of advanced features

Advanced Feature	Provides this functionality	And enters the license evaluation period when
Multi-Node Virtualization and Clustering	Clustering multiple storage systems to create pools of storage	You add two or more storage systems to a cluster in a management group.
Remote Copy	Creating secondary volumes and snapshots in remote locations	You create a remote volume in preparation for making a remote snapshot.
Managed Snapshot	Creating schedules to snapshot volumes	You create a schedule to snapshot a volume.
Multi-Site SAN	Multi-site clusters which synchronously and automatically mirror data between sites	You create a cluster with multiple sites.

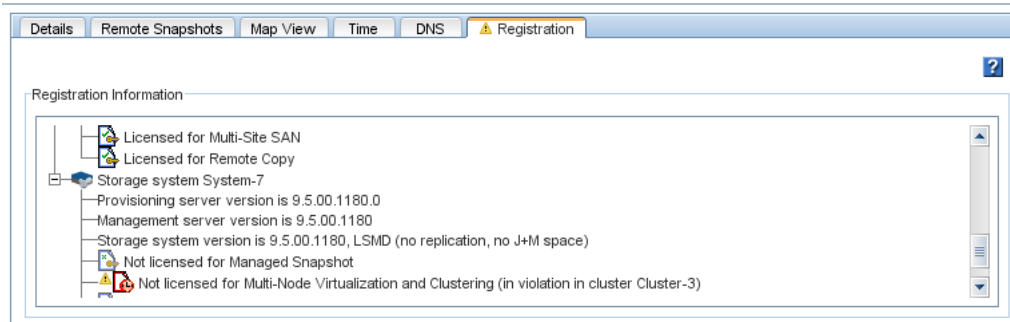
Tracking the time remaining in the evaluation period

Track the time left on your evaluation period by using either the management group Registration tab or the reminder notices that open periodically.

Identifying licensing status

You can check the status of licensing on individual advanced features by the icons displayed. The violation icon appears throughout the evaluation period.

Figure 105 Identifying the license status for advanced features



Backing out of Remote Copy evaluation

If you decide not to use Remote Copy and you have not obtained license keys by the end of the evaluation period, you must delete any remote volumes and snapshots you have configured. However, you can save the data in the remote snapshots before you delete them.

1. First, back up any volumes you plan to retain.
2. Next, safely back out of the Remote Copy evaluation as described in [Table 63 \(page 223\)](#), according to how you want to handle the data.

Table 63 Safely backing out of Remote Copy evaluation

Fate of Data in Remote Snapshots	Steps to Back Out
Removing data from the remote target	<ul style="list-style-type: none">• Delete the remote snapshots.• Delete the remote volume.
Retaining the data on the remote target	<ul style="list-style-type: none">• Make the remote volume into a primary volume.• Disassociate the primary and remote management groups, if the remote copy was between management groups.

Scripting evaluation

Application-based scripting is available for volume and snapshot features. You can create scripts to:

- Create snapshots
- Create remote volumes and snapshots

Because using scripts with advanced features starts the evaluation period without requiring that you use the CMC, you must first verify that you are aware of starting the evaluation clock when using scripting. If you do not enable the scripting evaluation period, any scripts you have running (licensed or not) will fail.

Turn on scripting evaluation

To use scripting while evaluating advanced features, enable the scripting evaluation period.

1. In the navigation window, select a management group.
2. Select the **Registration** tab.
3. Click **Registration Tasks**, and select **Feature Registration** from the menu.
4. Select the **Scripting Evaluation** tab.

5. Read the text, and select the box to enable the use of scripts during a license evaluation period.
6. Click **OK**.

Turn off scripting evaluation

Turn off the scripting evaluation period when you take either one of these actions:

- You obtain license keys for the feature you were evaluating.
- You complete the evaluation and decide not to license any advanced features.

Turning off the scripting evaluation ensures that no scripts continue to push the evaluation clock.

To turn off the scripting evaluation:

1. Select the management group.
2. Select the **Registration** tab.
3. Click **Registration Tasks**, and select **Feature Registration** from the menu.
4. Select the **Scripting Evaluation** tab.
5. Clear the check box.
6. Click **OK**.

Table 64 (page 224) describes additional steps to safely back out of the scripting evaluation.

Table 64 Safely backing out of scripting evaluation

Feature Being Evaluated	Steps to Back Out
Remote copy volumes and snapshots	<ul style="list-style-type: none">• Back out of any remote copy operation.• Delete any scripts.• Delete any primary or remote snapshots created by the scripts, as indicated by viewing Created By: on the snapshot Details tab.

Registering advanced features

When registering storage systems for advanced features, you must have your license entitlement certificate and submit to the website the appropriate storage system feature key(s) to obtain the license key(s). You then receive the license key(s) and apply them to the storage system(s).

Using license keys

License keys are assigned to individual storage systems. License keys can be added to storage systems either when they are in the Available Systems pool or after they are in a management group. One license key is issued per storage system, and that key licenses all the advanced features requested for that storage system.

NOTE: If you remove a storage system from a management group, the license key remains with that storage system. See for more information about removing storage systems from a management group.

Registering available storage systems for license keys

Storage systems that are in the Available Systems pool are licensed individually. You license an individual storage system on the Feature Registration tab for that system.

The Feature Registration tab displays the following information:

- The storage system feature key, used to obtain a license key
- The license key for that storage system, if one has been obtained
- The license status of all the advanced features

Submitting storage system feature keys

1. In the navigation window, select the storage system from the Available Systems pool for which you want to register advanced features.
2. Select the **Feature Registration** tab.
3. Select the Feature Key.
4. Right-click, and select **Copy**.
5. Use **Ctrl+V** to paste the feature key into a text editing program, such as Notepad.
6. Register and generate the license key at the Webware website:
<https://webware.hp.com>

Entering license keys to storage systems

When you receive the license keys, add them to the storage systems.

1. In the navigation window, select the storage system from the Available Systems pool.
2. Select the **Feature Registration** tab.
3. Click **Feature Registration Tasks**, and select **Edit License Key** from the menu.
4. Copy and paste the Feature Key into the Edit License Key window.

NOTE: When you paste the license key into the window, be sure there are no leading or trailing spaces in the box. Such spaces prevent the license key from being recognized.

5. Click **OK**.

The license key appears in the Feature Registration window.

Figure 106 Storage system with a license key



Registering storage systems in a management group

Storage systems that are in a management group can be licensed through the management group. License the storage systems on the Registration tab for the management group.

The Registration tab displays the following information:


- The license status of all the advanced features, including the progress of the evaluation period and which advanced features are in use and not licensed
- Version information about software components of the operating system
- Customer information (optional)

Submitting storage system feature keys

Submit the feature keys of all the storage systems in the management group.

1. In the navigation window, select the management group for which you want to register advanced features.
2. Click the **Registration** tab.

Figure 107 Registering advanced features for a management group



3. Click **Registration Tasks**, and select **Feature Registration** from the menu.
The Feature Registration window lists all the storage systems in that management group.

Figure 108 Selecting the feature key



4. For each storage system listed in the window, select the Feature Key.
5. Press **Ctrl+C** to copy the Feature Key.
6. Use **Ctrl+V** to paste the feature key into a text editing program, such as Notepad.
7. Go to the HP Webware site to register and generate the license key.
<https://webware.hp.com>

NOTE: Record the host name or IP address of the storage system along with the feature key. This record will make it easier to add the license key to the correct storage system when you receive it.

Entering license keys

When you receive the license keys, add them to the storage systems in the Feature Registration window.

1. In the navigation window, select the management group.
2. Select the **Registration** tab.
3. Click **Registration Tasks**, and select **Feature Registration** from the menu.
4. Do one of the following.

To enter a license key for one storage system in the management group	To enter license keys for multiple storage systems in the management group
<p>a. Select a storage system, and click Edit License Key.</p> <p>b. Copy and paste the appropriate license key for that storage system into the window.</p> <p>NOTE: When you cut and paste the license key into the window, ensure that there are no leading or trailing spaces in the box. Such spaces prevent the license key from being recognized.</p> <p>c. Click OK.</p> <p>The license key appears in the Feature Registration window.</p> <p>d. Click OK again to exit the Feature Registration window.</p>	<p>a. Click Import License Keys.</p> <p>b. Click Browse to navigate to the license key files you downloaded from https://webware.hp.com.</p> <p>c. Select each .dat file that corresponds to the storage systems in the management group, and click Open.</p> <p>Each license key is in a separate .dat file. The file name has the storage system's feature key in the file name, as follows: xxxxxxxxxxxxx_xxxxxxx_AA.BB.CC.DD.EE.FF_x.dat. Be sure the AA.BB.CC.DD.EE.FF part of each file name matches the feature key of a storage system. If an error message appears, the error text describes the problem.</p> <p>d. Click Apply License Keys.</p> <p>e. Check the Import Summary window for any errors.</p> <p>The error messages explain the problem. A green check mark shows the license key has been applied.</p> <p>f. Click Close.</p> <p>If the import process had errors, you return to the Import License Key window. Click Browse to select a different license key file. If you click Cancel to exit the window, any license keys with a green check mark were applied. If the import process completed successfully, click Close to exit the window.</p>

Saving license key information

For record-keeping, save the license information to a .txt file when you have entered all the license keys.

1. Click **Registration Tasks** on the management group Registration tab.
2. Select **Save Information to File** from the menu.
3. Navigate to the location where you want to save the license key information.
4. Enter a name for the registration information file, and click **Save**.

Saving and editing your customer information

Save your customer profile, registrations, and licensing information as a text file. Then, if you lose a storage system, the information can help you rebuild a new storage system.

Make a customer information file for each management group in your SAN.

- Create or edit your customer profile.
- Save the customer profile to a computer that is not part of your SAN.

Editing your customer information file

Occasionally, you may want to change some of the information in your customer profile. For example, if your company moves, or contact information changes.

1. In the navigation window, select a management group.
2. Click the **Registration** tab.
3. Click **Registration Tasks**, and select **Edit Customer Information** from the menu.
4. Fill in or change any of the information on this window.
5. Click **OK** when you are finished.

Saving your customer information

Be sure you have filled in the customer profile window correctly before saving this file. In addition to the customer information, the file you save contains registration and licence key information.

Save a customer information file for each management group in your storage system.

1. In the navigation window, select a management group.
2. Click the **Registration** tab.
3. Click **Registration Tasks**, and select **Save Information to File** from the menu.
4. In the Save window, navigate to the directory where you want to save the license key and customer information file.
5. In the File Name field, enter a name for the file, which defaults to a .txt file.
6. Click **Save**.

Verify the information by viewing the saved .txt file.

19 iSCSI and the HP P4000 SAN Solution

The SAN/iQ software uses the iSCSI protocol to let servers access volumes. For fault tolerance and improved performance, use a VIP and iSCSI load balancing when configuring server access to volumes.

The following concepts are important when setting up clusters and servers in the SAN/iQ software:

- [“Virtual IP addresses” \(page 229\)](#)
- [“iSNS server” \(page 229\)](#)
- [“iSCSI load balancing” \(page 230\)](#)
- [“Authentication \(CHAP\)” \(page 230\)](#)
- [“iSCSI and CHAP terminology” \(page 231\)](#)
- [“About HP DSM for MPIO” \(page 234\)](#)

Number of iSCSI sessions

For information about the recommended maximum number of iSCSI sessions that can be created in a management group, see [“Configuration Summary overview” \(page 108\)](#).

Virtual IP addresses

A virtual IP (VIP) address is a highly available IP address which ensures that if a storage system in a cluster becomes unavailable, servers can still access a volume through the other storage systems in the cluster.

Your servers use the VIP to discover volumes on the SAN. The SAN uses the iqn from the iSCSI initiator to associate volumes with the server.

A VIP is required for a fault tolerant iSCSI cluster configuration, using VIP load balancing or the HP DSM for MPIO.

When using a VIP, one storage system in the cluster hosts the VIP. All I/O goes through the VIP host. You can determine which storage system hosts the VIP by selecting the cluster, then clicking the iSCSI tab.

Requirements for using a virtual IP address

- For standard clusters (not multi-site clusters), storage systems occupying the same cluster must be in the same subnet address range as the VIP.
- The VIP must be routable regardless of which storage system it is assigned to.
- iSCSI servers must be able to ping the VIP when it is enabled in a cluster.
- The VIP address must be different than any storage system IPs on the network.
- The VIP address must be a static IP address reserved for this purpose.
- All iSCSI initiators must be configured to connect to the VIP address for the iSCSI failover to work properly.

iSNS server

An iSNS server simplifies the discovery of iSCSI targets on multiple clusters on a network. If you use an iSNS server, configure your cluster to register the iSCSI target with the iSNS server. You can use up to 3 iSNS servers, but none are required.

iSCSI load balancing

Use iSCSI load balancing to improve iSCSI performance and scalability by distributing iSCSI sessions for different volumes evenly across storage systems in a cluster. iSCSI load balancing uses iSCSI Login-Redirect. Only initiators that support Login-Redirect should be used.

When using VIP and load balancing, one iSCSI session acts as the gateway session. All I/O goes through this iSCSI session. You can determine which iSCSI session is the gateway by selecting the cluster, then clicking the iSCSI Sessions tab. The Gateway Connection column displays the IP address of the storage system hosting the load balancing iSCSI session.

Configure iSCSI load balancing when setting up servers. See “iSCSI and the HP P4000 SAN Solution” (page 229).

Requirements

- Cluster configured with a virtual IP address. See “Virtual IP addresses” (page 229).
- A compliant iSCSI initiator.

Compliant iSCSI initiators

A compliant initiator supports iSCSI Login-Redirect AND has passed HP’ test criteria for iSCSI failover in a load balanced configuration.

Find information about which iSCSI initiators are compliant by clicking the link in the New or Edit Server window.

The link opens to the <http://www.hp.com/go/P4000compatibility> where you log in and search for the *HP P4000 SAN Solutions Compatibility Matrix*.

If your initiator is not on the list, do not enable load balancing.

Authentication (CHAP)

Server access with iSCSI can use the following authentication methods:

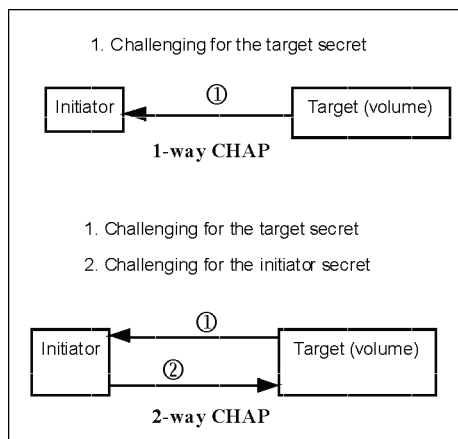
- Initiator node name (single host)
- CHAP (Challenge-Handshake Authentication Protocol), which can support single or multiple hosts.

NOTE: The iSCSI terminology in this discussion is based on the Microsoft iSCSI Initiator terminology. For the terms used in other common operating systems, see “iSCSI and CHAP terminology” (page 231).

CHAP is a standard authentication protocol. The SAN/iQ software supports the following configurations:

- No CHAP—Authorized initiators can log in to the volume without proving their identity. The target does not challenge the server.
- 1-way CHAP—Initiators must log in with a target secret to access the volume. This secret proves the identity of the initiator to the target.
- 2-way CHAP—Initiators must log in with a target secret to access the volume as in 1-way CHAP. In addition, the target must prove its identity to the initiator using the initiator secret. This second step prevents target spoofing.

Figure 109 Differentiating types of CHAP



CHAP is optional. However, if you configure 1-way or 2-way CHAP, you must remember to configure both the server and the iSCSI initiator with the appropriate characteristics. [Table 65 \(page 231\)](#) lists the requirements for configuring CHAP.

Requirements for configuring CHAP

Table 65 Configuring iSCSI CHAP

CHAP Level	What to Configure for the Server in the SAN/iQ Software	What to Configure in the iSCSI Initiator
CHAP not required	Initiator node name only	No configuration requirements
1-way CHAP	<ul style="list-style-type: none">• CHAP name*• Target secret	Enter the target secret (12-character minimum) when logging on to available target.
2-way CHAP	<ul style="list-style-type: none">• CHAP name*• Target secret• Initiator secret	<ul style="list-style-type: none">• Enter the initiator secret (12-character minimum).• Enter the target secret (12-character minimum).

* If using CHAP with a single node only, use the initiator node name as the CHAP name.

iSCSI and CHAP terminology

The iSCSI and CHAP terms used vary based on the operating system and iSCSI initiator you are using. The table below lists the terms for two common iSCSI initiators.

Table 66 iSCSI terminology

SAN/iQ CMC	Microsoft	VMWare	Linux
Initiator Node Name	Initiator Node Name	iSCSI Name	Refer to the documentation for the iSCSI initiator you are using. Linux iSCSI initiators may use a command line interface or a configuration file.
CHAP Name	Not used	CHAP Name	
Target Secret	Target Secret	CHAP Secret	
Initiator Secret	Secret	N/A	

NOTE: The initiator node name and secrets set in the SAN/iQ CMC must match what you enter in the server's iSCSI initiator exactly.

Sample iSCSI configurations

Figure 110 (page 232) illustrates the configuration for a single host authentication with CHAP not required with Microsoft iSCSI.

Figure 110 Viewing the MS iSCSI initiator to copy the initiator node name

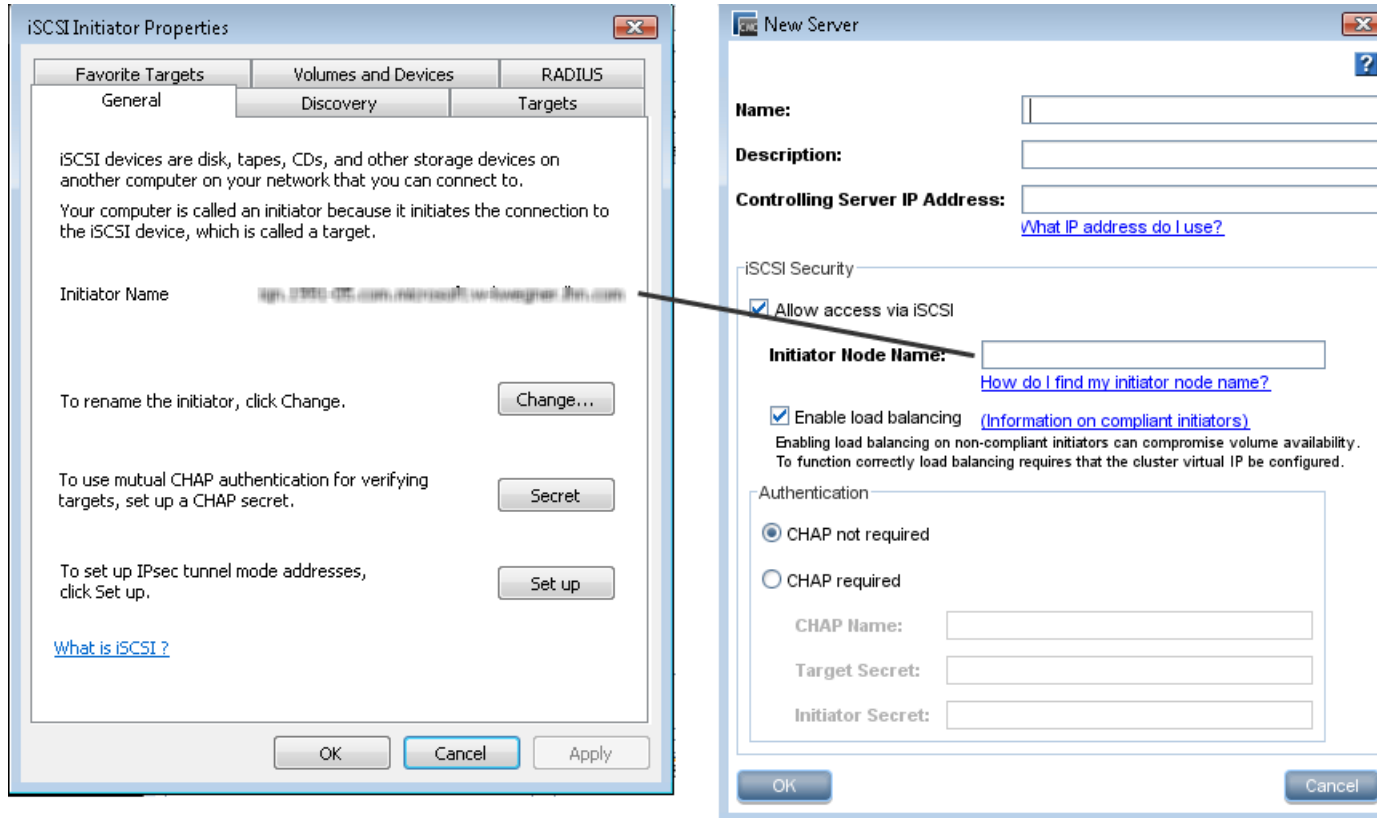


Figure 111 (page 233) illustrates the configuration for a single host authentication with 1-way CHAP required.

Figure 111 Configuring iSCSI (shown in the MS iSCSI initiator) for a single host with CHAP

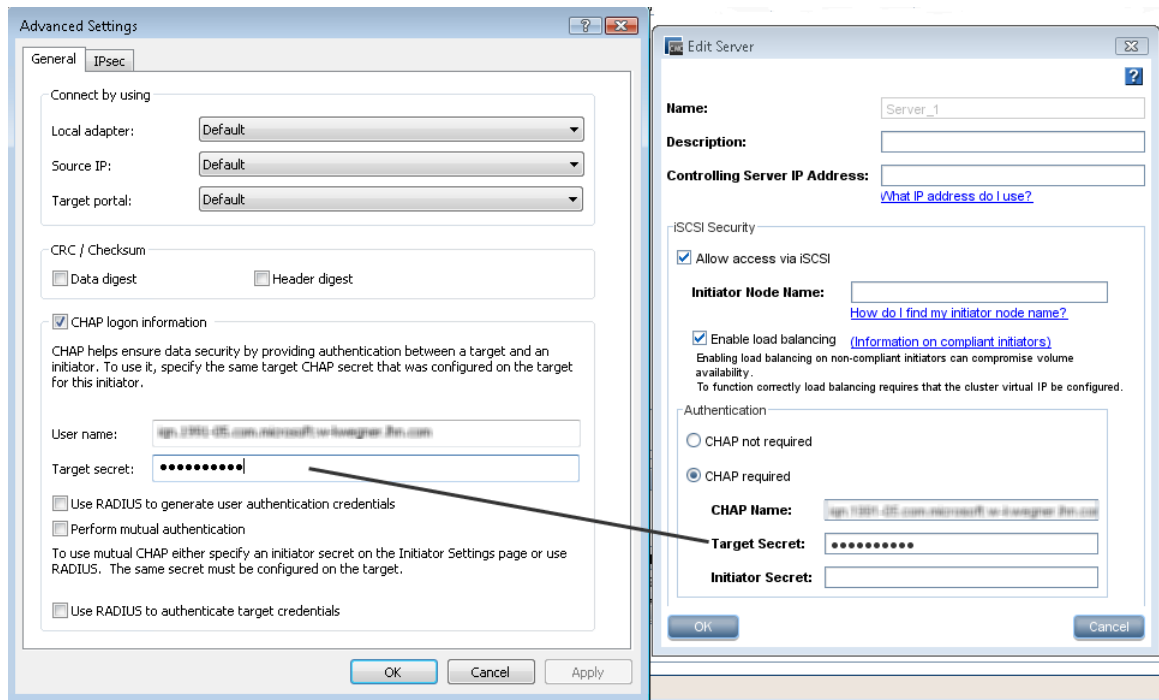
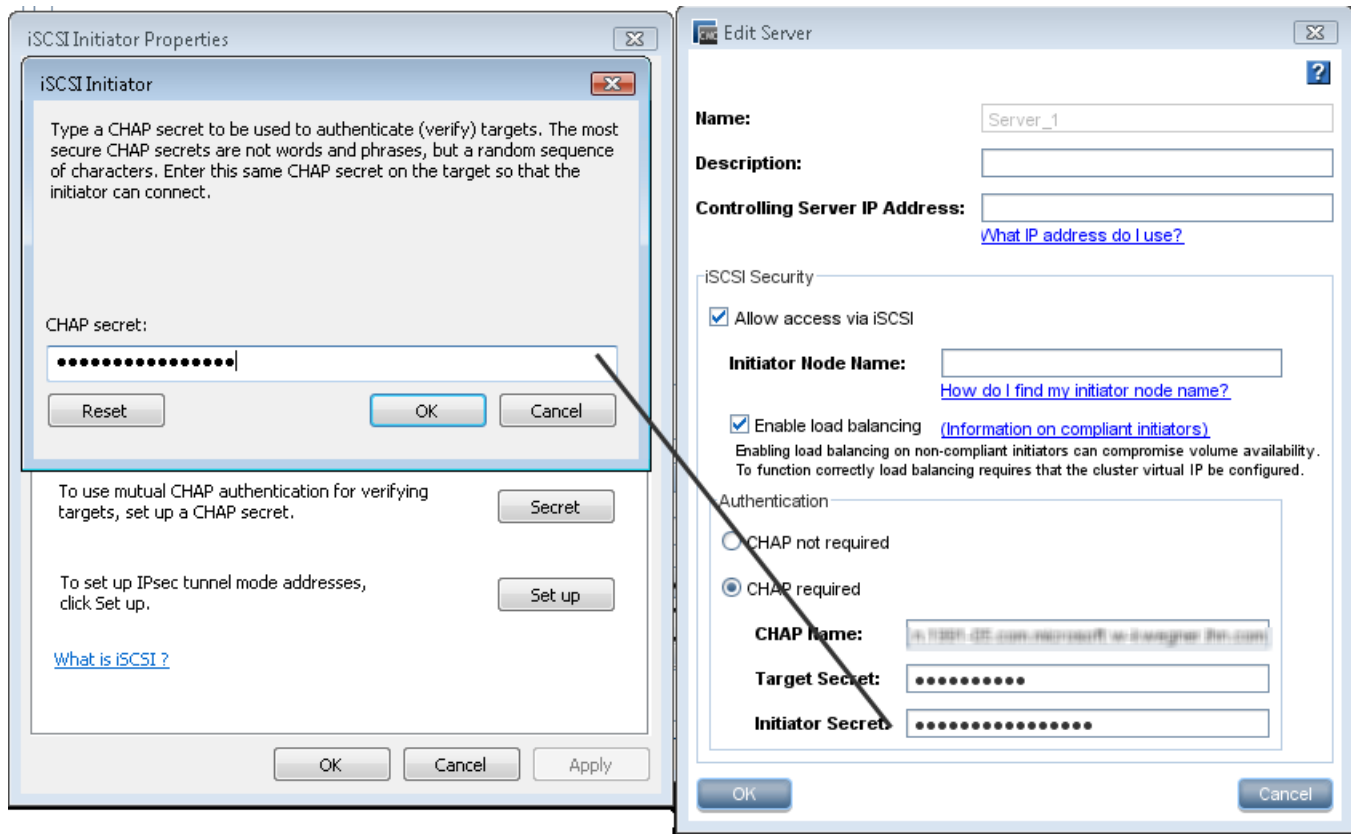


Figure 112 (page 234) illustrates the configuration for a single host authentication with 2-way CHAP required.

Figure 112 Adding an initiator secret for 2-way CHAP (shown in the MS iSCSI initiator)



CAUTION: Without the use of shared storage access (host clustering or clustered file system) technology, allowing more than one iSCSI application server to connect to a volume concurrently without cluster-aware applications and/or file systems in read/write mode could result in data corruption.

NOTE: If you enable CHAP on a server, it will apply to all volumes for that server.

Best practice

In the Microsoft iSCSI Initiator, target and initiator secrets are not displayed. Keep a separate record of the iSCSI Initiator CHAP information and the corresponding server information.

About HP DSM for MPIO

If you are using the HP DSM for MPIO, you can use HP DSM for MPIO to access volumes. For more information about HP DSM for MPIO, refer to the HP P4000 DSM for MPIO Deployment Guide.

You can see if you are using HP DSM for MPIO in the CMC by selecting a server in a management group, then clicking the Volumes and Snapshots tab. The Gateway Connection column shows multiple connections labeled as DSM.

When accessing volumes from a server using HP DSM for MPIO, keep in mind the following:

- HP DSM for MPIO and the Microsoft MPIO must be installed on the server.
- With these installed, servers automatically use HP DSM for MPIO when you log on to volumes from the iSCSI initiator.
- If you have dual storage NICs in your server, you can select the Enable multi-path option when logging on to the volume, and log on from each NIC.

20 Using the Configuration Interface

The Configuration Interface is the command line interface that uses a direct connection with the storage system.

You may need to access the Configuration Interface if all network connections to the storage system are disabled. Use the Configuration Interface to perform the following tasks.

- Add storage system administrators and change passwords
- Access and configure network interfaces
- Delete a NIC bond
- Set the TCP speed and duplex, or edit the frame size
- Remove the storage system from a management group
- Reset the storage system configuration to factory defaults

Connecting to the Configuration Interface

Accessing the Configuration Interface is accomplished by

- Attaching a keyboard and monitor (KVM) to the storage system serial port (preferred) or
- Attaching a PC or a laptop using a null modem cable and connecting to the Configuration Interface with a terminal emulation program.

Establishing a terminal emulation session on a Windows system

On the PC or laptop attached directly to the storage system with a null modem cable, open a session with a terminal emulation program such as HyperTerminal or ProComm Plus.

Use the following settings: 19200, 8-N-1

When the session is established, the Configuration Interface window opens.

Establishing a terminal emulation session on a Linux/UNIX system

If using Linux, create the following configuration file. You must create the file as root, or root must change permissions for /dev/cua0 in order to create the config file in /etc/.

1. Create the /etc/minirc.NSM with the following parameters:

```
# Begin HP LeftHand Networks NSM configuration
# Machine-generated file - use "minicom -s" to
# change parameters
pr port = /dev/cua0
pu baudrate = 19200
pu bits = 8
pu parity = N
pu stopbits = 1
pu autobaud = Yes
pu backspace = DEL
pu hasdcd = No
pu rtscts = No
pu xonxoff = Yes
pu askndndir = Yes
# End HP LeftHand Networks NSM configuration
```
2. Start xterm as follows:

```
$ xterm
```

3. In the xterm window, start minicom as follows:

```
$ minicom -c on -l NSM
```

Opening the Configuration Interface from the terminal emulation session

1. Press **Enter** when the terminal emulation session is established.
2. Enter `start`, and press **Enter** at the log in prompt.
3. When the session is connected to the storage system, the Configuration Interface window opens.

Logging in to the Configuration Interface

Once you have established a connection to the storage system, log in to the Configuration Interface.

Table 67 Logging in depends on where the storage system is

If the storage system is in	From Configuration Interface entry window
Available Systems pool	Press Enter to log in. The Configuration Interface main menu opens.
Management group	<ol style="list-style-type: none">1. Press Enter to log in.2. Enter the user name and password of the administrative user created for the management group.3. Tab to Login, and press Enter.

NOTE: This user is viewable in the CMC under the management group Administration category.

Configuring administrative users

Use the Configuration Interface to add new administrative users or to change administrative passwords. You can only change the password for the administrative user that you used to log in to the Configuration Interface.

1. On the Configuration Interface main menu, tab to **General Settings**, and press **Enter**.
2. To add an administrative user, tab to **Add Administrator**, and press **Enter**. Then enter the new user's name and password. Confirm password, tab to **OK**, and press **Enter**.
3. To change the password for the user that you are currently logged in as, tab to **Change Password**, and press **Enter**. Then enter the new password. Confirm password, tab to **OK**, and press **Enter**.
4. On the General window, tab to **Done**, and press **Enter**.

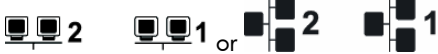
Configuring a network connection

The storage system comes with two Ethernet interfaces. [Table 68 \(page 236\)](#) lists where interfaces are labeled and the label name.

Table 68 Identifying Ethernet interfaces on the storage system

Ethernet Interfaces	
Where labeled	What the label says
TCP/IP Network configuration category in the CMC	Name:
<ul style="list-style-type: none">• TCP/IP tab• TCP Status tab	eth0, eth1 Motherboard:Port0, Motherboard:Port1 G4-Motherboard:Port1, G4-Motherboard:Port2

Table 68 Identifying Ethernet interfaces on the storage system *(continued)*

Ethernet Interfaces	
Where labeled	What the label says
	Motherboard:Port1, Motherboard:Port2
Configuration Interface	Intel Gigabit Ethernet or Broadcom Gigabit Ethernet
Label on the back of the storage system	Eth0, Eth1, or a graphical symbol similar to the following: 

Once you have established a connection to the storage system using a terminal emulation program, you can configure an interface connection using the Configuration Interface.

1. On the Configuration Interface main menu, tab to **Network TCP/IP Settings**, and press **Enter**.
2. Tab to select the network interface that you want to configure, and press **Enter**.
3. Enter the host name, and tab to the next section to configure the network settings.

NOTE: If you specify an IP address, Gateway is a required field. If you do not have a Gateway, enter 0.0.0.0 for the Gateway address.

4. Tab to **OK**, and press **Enter** to complete the network configuration.
5. Press **Enter** on the confirmation window.
6. Open the CMC, and locate the storage system using the Find function.

Deleting a NIC bond

You can delete the following NIC bonds using the Configuration Interface:

- Active-Passive bond
- Link Aggregation Dynamic Mode bond
- Adaptive Load Balancing bond

For more information about creating and configuring NIC bonds, see [“Configuring network interface bonds” \(page 55\)](#).

When you delete an Active-Passive bond, the primary interface assumes the IP address and configuration of the deleted logical interface. The other NIC is disabled and its IP address is set to 0.0.0.0.

When you delete a Link Aggregation Dynamic Mode or an Adaptive Load Balancing bond, eth0 or motherboard: port 1 retains the IP address of the deleted logical interface. The other NIC is disabled, and its IP address is set to 0.0.0.0.

1. On the Configuration Interface main menu, tab to **Network TCP/IP**, Settings and press **Enter**.
In the Available Network Devices window that opens, the logical bond is the only interface listed.
2. Tab to select the bond, and press **Enter**.
3. Tab to **Delete Bond**, and press **Enter**.
4. Press **Enter** on the confirmation window.
5. On the Available Network Devices window, tab to **Back**, and press **Enter**.

Setting the TCP speed, duplex, and frame size

You can use the Configuration Interface to set the TCP speed, duplex, and frame size of a network interface.

TCP speed and duplex. You can change the speed and duplex of an interface. If you change these settings, you must ensure that both sides of the NIC cable are configured in the same manner. For example, if the storage system is set for Auto/Auto, the switch must be set the same. For more information about TCP speed and duplex settings, see [“Managing settings on network interfaces” \(page 50\)](#).

Frame size. The frame size specifies the size of data packets that are transferred over the network. The default Ethernet standard frame size is 1500 bytes. The maximum allowed frame size is 9000 bytes.

Increasing the frame size improves data transfer speed by allowing larger packets to be transferred over the network and by decreasing the CPU processing time required to transfer data. However, increasing the frame size requires that routers, switches, and other devices on your network support that frame size.

For more information about setting a frame size that corresponds to the frame size used by routers, switches, and other devices on your network, see [“Changing NIC frame size” \(page 51\)](#).

1. On the Configuration Interface main menu, tab to **Network TCP Status** and press **Enter**.
2. Tab to select the network interface for which you want to set the TCP speed and duplex, and press **Enter**.
3. To change the speed and duplex of an interface, tab to a setting in the Speed / Duplex list.
4. To change the frame size, select **Set To in the Frame Size** list. Then tab to the field to the right of Set To, and enter a frame size.

The frame size value must be between 1500 bytes and 9000 bytes.

5. On the Network TCP Status window, tab to **OK**, and press **Enter**.
6. On the Available Network Devices window, tab to **Back**, and press **Enter**.

Removing a storage system from a management group

-
- △ **CAUTION:** Removing a storage system from a management group deletes all data from the storage system, clears all information about the management group, and reboots the storage system.
-

1. On the Configuration Interface main menu, tab to **Config Management**, and press **Enter**.
2. Tab to **Remove from management group**, and press **Enter**.
3. Tab to **Ok**, and press **Enter** to confirm.
4. On the Configuration Management window, tab to **Done**, and press **Enter**.

Resetting the storage system to factory defaults

-
- △ **CAUTION:** Resetting the storage system to factory defaults deletes all data and erases the configuration of the storage system, including administrative users and network settings.
-

1. On the Configuration Interface main menu, tab to **Config Management**, and press **Enter**.
2. Tab to **Reset to factory defaults**, and press **Enter**.
3. Tab to **Ok**, and press **Enter** to confirm the reset.
4. On the Configuration Management window, tab to **Done**, and press **Enter**.

21 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/e-updates>

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

HP Insight Remote Support Software

HP strongly recommends that you install HP Insight Remote Support software to complete the installation or upgrade of your product and to enable enhanced delivery of your HP Warranty, HP Care Pack Service or HP contractual support agreement. HP Insight Remote Support supplements your monitoring, 24x7 to ensure maximum system availability by providing intelligent event diagnosis, and automatic, secure submission of hardware event notifications to HP, which will initiate a fast and accurate resolution, based on your product's service level. Notifications may be sent to your authorized HP Channel Partner for on-site service, if configured and available in your country. The software is available in two variants:

- **HP Insight Remote Support Standard:** This software supports server and storage devices and is optimized for environments with 1-50 servers. Ideal for customers who can benefit from proactive notification, but do not need proactive service delivery and integration with a management platform.
- **HP Insight Remote Support Advanced:** This software provides comprehensive remote monitoring and proactive service support for nearly all HP servers, storage, network, and SAN environments, plus selected non-HP servers that have a support obligation with HP. It is integrated with HP Systems Insight Manager. A dedicated server is recommended to host both HP Systems Insight Manager and HP Insight Remote Support Advanced.

Details for both versions are available at: <http://h18004.www1.hp.com/products/servers/management/insight-remote-support/overview.html>

To download the software, go to Software Depot:

<https://h20392.www2.hp.com/portal/swdepot/index.do>.

Select Insight Remote Support from the menu on the right.

New and changed information in this edition

The following additions and changes have been made for this edition:

- The following information has been updated:
 - P4000 SAN Solution and user documentation have been rebranded
 - New SAN Status Home page is added to the CMC, providing a graphic view for monitoring the SAN
 - New Best Practice analyzers have been added
 - Management Groups, Clusters and Volume wizard has been updated with elements to ensure a highly available SAN configuration
 - New snapshot functionality to quiesce application servers on VMware Servers before taking snapshots

Related information

The following documents [and websites] provide related information:

You can find the documents referenced in this guide from the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

In the Storage section, select **Disk Storage Systems**→**P4000 SAN Solutions**→**HP P4000 G2 SAN Solutions**.

HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- http://www.hp.com/service_locator
- <http://www.hp.com/support/manuals>
- <http://www.hp.com/support/downloads>
- <http://www.hp.com/storage/whitepapers>

Customer self repair

HP customer self repair (CSR) programs allow you to repair your P4000 product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider, or see the CSR website:

<http://www.hp.com/go/selfrepair>

A Replacing disks reference

NOTE: RAID refers to the disk level RAID set on an individual storage system. Network RAID refers to the data protection level set on volumes when they are created. The data protection level is always referred to as Network RAID. Disk RAID is always referred to as RAID.

This chapter describes the disk replacement procedures for cases in which you do not know which disk to replace and/or you must rebuild RAID on the entire storage system. For example, if RAID has gone off unexpectedly, you need Customer Support to help determine the cause, and if it is a disk failure, to identify which disk must be replaced.

Replacing disks and rebuilding data

Single disk replacements in storage systems where RAID is running, but may be degraded, can be accomplished by following the procedures described in [“Replacing disks” \(page 242\)](#).

The following situations may require consulting with Customer Support to identify bad disks, and then following the procedures in this chapter to rebuild the data (when configured for data protection) on the storage system:

- RAID 0 (Stripe) — RAID is off due to a failed disk.
- RAID 5, 5+spare (Stripe with parity), and 50 — If multiple disks need to be replaced, they must be identified and replaced, and the data on the entire storage system rebuilt.
- RAID 10/1+0 (Mirror and Stripe) — Can sustain multiple disk replacements. However, Customer Support must identify if any two disks are from the same mirror set, and then the data on the entire storage system needs to be rebuilt.
- RAID 6 (Stripe with dual parity) — If multiple disks need to be replaced, they must be identified and replaced, and the data on the entire storage system rebuilt.

Before you begin

1. Know the name and physical location of the storage system that needs the disk replacement.
 2. Know the physical position of the disk in the storage system.
 3. Have the replacement disk ready, and confirm that it is the right size and has the right carrier.
- For confirmation on which disks need to be replaced, contact customer support.

Prerequisites

- All Network RAID-10, Network RAID-10+1, Network RAID-10+2, Network RAID-5, and Network RAID-6 volumes and snapshots should show a status of Normal. Network RAID-0 volumes may be offline.
- If volumes or snapshots are Network RAID-0 and online, change them to Network RAID-10 or other Network RAID level before replacing the disk.
- If the cluster does not have enough space for the replication, take a backup of the volumes or snapshots, and then delete them from the cluster. After the disk replacement is complete, recreate the volumes, and restore the data from the backup.
- Any volumes or snapshots that were being deleted should have finished deleting.
- Write down the order in which the storage systems are listed in the Edit Cluster window. You must ensure that they are all returned to that order when the repair is completed.

Replacing disks

Use this procedure when any of the following occurs.

- RAID on a storage system configured with RAID 0 goes off because of a disk failure.
- Multiple disks need to be replaced on a storage system with RAID 5, RAID 50, or RAID 6.
- Multiple disks on the same mirror set need to be replaced on a storage system with RAID 10.

Verify storage system not running a manager

Verify that the storage system that needs the disk replacement is not running a manager.

1. Log in to the management group.
2. Select the storage system in the navigation window, and review the Details tab information. If the Storage System Status shows Manager Normal, and the Management Group Manager shows Normal, then a manager is running and needs to be stopped.

Stopping a manager

1. To stop a manager, right-click the storage system in the navigation window, and select **Stop Manager**.

When the process completes successfully, the manager is removed from the Status line in the Storage System box, and the Manager changes to No in the Management Group box.

2. If you stop a manager, the cluster will be left with an even number of managers. To ensure that the cluster has an odd number of managers, do one of these tasks:
 - Start a manager on another storage system.
 - Add a virtual manager to the management group by right-clicking on the management group name in the navigation window and selecting **Add Virtual Manager**.

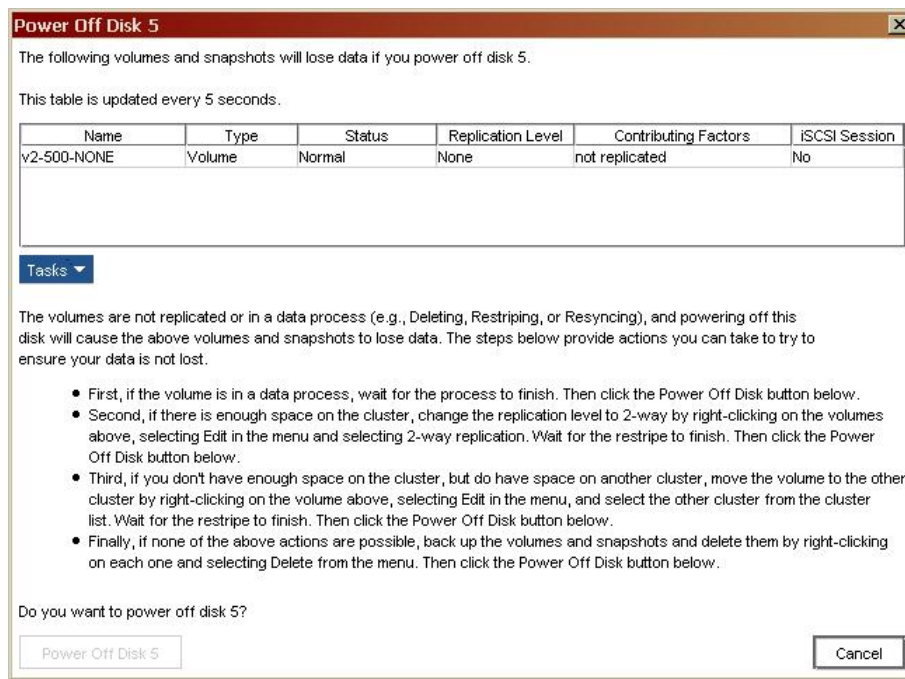
Repair the storage system

Use the Repair Storage System function to replace a disk and trigger only one resync of the data.

Prerequisite

If there are Network RAID-0 volumes that are offline, you must either replicate them or delete them before you can proceed with this step. You see the message shown in [Figure 113 \(page 243\)](#) in this case.

Figure 113 Warning if volumes are Network RAID-0



Right-click the storage system in the navigation window, and select **Repair Storage System**. A “ghost” image takes the place of the storage system in the cluster, with the IP address serving as a place holder. The storage system itself moves from the management group to the Available Systems pool.

NOTE: If the storage system does not appear in the Available Systems pool, use the Find menu option to relocate it.

Replace the disk

In the DL320s (NSM 2120), HP LeftHand P4300, HP LeftHand P4500

For the DL320s (NSM 2120), HP LeftHand P4300, and HP LeftHand P4500, use the disk replacement procedures in “Replacing a disk in a hot-swap storage system ” (page 46).

Rebuilding data

The following steps take you through the steps to first rebuild RAID on the storage system and then to rebuild data on the storage system after it is added to the management group and cluster.

Recreate the RAID array

1. Select the **Storage** category, and select the **RAID Setup** tab.
2. Click **RAID Setup Tasks**, and select **Reconfigure RAID**.

The RAID Status changes from Off to Normal.

NOTE: If RAID reconfigure reports an error, reboot the storage system, and try reconfiguring the RAID again. If this second attempt is not successful, call customer support.

Checking progress for RAID array to rebuild

For DL320s (NSM 2120), HP P4500 and P4500 G2, HP P4300 and P4300 G2.

Use the Hardware Information report to check the status of the RAID rebuild.

1. Select the **Diagnostics** category, and select the **Hardware Information** tab.
2. Click the **Click to Refresh** link, and scroll down to the RAID section of the Hardware report, shown in Figure 114 (page 244).
You can view the RAID rebuild rate and percent complete.
3. Click **Hardware Information Tasks**, and select **Refresh** to monitor the ongoing progress.

Figure 114 Checking RAID rebuild status

Last Refreshed: 7/9/10 9:17:16 AM MDT		
Item	Value	
Drive Status	Status	Health
Drive 1	Active	Normal
Drive 2	Active	Normal
Drive 3	Active	Normal
Drive 4	Active	Normal
Drive 5	Active	Normal
Drive 6	Active	Normal
Drive 7	Active	Normal
Drive 8	Rebuilding	Normal
Drive 9	Active	Normal
Drive 10	Active	Normal
Drive 11	Active	Normal
Drive 12	Active	Normal
RAID	Rebuilding	
Rebuild Rate	Low	
Unused Devices	Disk 8_Rebuilding	
Statistics	2 Units	
Unit 1	/dev/cciss/c0d1 : DATA Partition Raid 5 4284 08 GB Normal	

Return storage system to cluster

Return the repaired storage system to the cluster.

1. In the navigation window, right-click the storage system, and select **Add to Existing Management Group**.
2. Select from the list the Group Name that the storage system used to belong to and click **Add**.
The storage system appears in the management group and the icon in the navigation window flashes for a few minutes as it initializes.

Restarting a manager

Before proceeding, make sure that the storage system has finished initializing and is completely added to the management group.

If necessary, ensure that after the repair you have the appropriate configuration of managers. If there was a manager running on the storage system before you began the repair process, you may start a manager on the repaired storage system as necessary to finish with the correct number of managers in the management group.

If you added a virtual manager to the management group, you must first delete the virtual manager before you can start a regular manager.

1. Right-click the virtual manager, and select **Stop Virtual Manager**.
2. Right-click the virtual manager, and select **Delete Virtual Manager**.
3. Right-click the storage system, and select **Start Manager**.

Add repaired system to cluster

1. After the initialization completes, right-click the cluster, and select **Edit Cluster**. The list of the storage systems in the cluster should include the ghost IP address.
You now need to add the repaired storage system to the cluster in the spot held by the ghost IP address.
2. Select the ghost storage system (the IP address in the list) and click **Exchange System**.
3. Select the repaired storage system to exchange for the ghost storage system and click **OK**.
The storage system returns to its original position in the cluster and volumes in the cluster proceed to resync.

Table 69 Replacing the ghost storage system with the repaired storage system

	Storage Systems in Cluster
Before rearranging	<ul style="list-style-type: none">• Storage system A• <IP Address>• Storage system C
After rearranging	<ul style="list-style-type: none">• Storage system A• Storage system B• Storage system C

NOTE: If you do not arrange the storage systems to match their original order, the data in the cluster is rebuilt across all the storage systems instead of just the repaired storage system. This total data rebuild takes longer to complete and increases the chance of a second failure during this period.

To ensure that only the repaired storage system goes through the rebuild, before you click the OK button in the Edit Cluster window, double-check that the order of the storage systems in the cluster list matches the original order.

Rebuild volume data

After the storage system is successfully added back to the cluster, the adjacent storage systems start rebuilding data on the repaired storage system.

1. Select the cluster, and select the **Disk Usage** tab.
2. Verify that the disk usage on the repaired storage system starts increasing.
3. Verify that the status of the volumes and snapshots is Restriping.

Depending on the usage, it may take anywhere from a few hours to a day for the data to be rebuilt on the repaired storage system.

Controlling server access

Use the Local Bandwidth Priority setting to control server access to data during the rebuild process.

- When the data is being rebuilt, the servers that are accessing the data on the volumes might experience slowness. Reduce the Local Bandwidth Priority to half of its current value for immediate results.
- Alternatively, if server access performance is not a concern, raise the Local Bandwidth Priority to increase the data rebuild speed.

Change local bandwidth priority

1. Right-click the management group, and select **Edit Management Group**.
The current Bandwidth Priority value indicates that each manager in that management group will use that much bandwidth to transfer data to the repaired storage system. Make a note of the current value so it can be restored after the data rebuild completes.
2. Change the bandwidth value as desired, and click **OK**.

Remove ghost storage system

Remove the ghost storage system after the data is rebuilt.

The data is rebuilt on the storage system when two conditions are met:

- The repaired storage system's disk usage matches the usage of the other storage systems in the cluster.
- The status of the volume and snapshots goes back to Normal.

The ghost IP address showing outside the cluster can now be removed from the management group.

1. Right-click the ghost IP, address and select **Remove from Management Group**.
2. If you have adjusted/reduced the Local Bandwidth Priority of the management group while the data was being rebuilt, change it back to the original value.

At this point, the disk(s) in the storage system are successfully replaced, the data will be fully rebuilt on that storage system, and the management group configuration (like number of managers, quorum, local bandwidth, and so on) will be restored to the original settings.

Finishing up

1. Contact Customer Support for an RA number.
2. Return the original disks for failure analysis using the prepaid packing slip in the replacement package. Put the RA number on the package as instructed by Customer Support.

B Third-party licenses

The software distributed to you by HP includes certain software packages indicated to be subject to one of the following open source software licenses: GNU General Public License ("GPL"), the GNU Lesser General Public License ("LGPL"), or the BSD License (each, an "OSS Package"). Refer to the license_readme file included with the software for additional information regarding the application of these licenses to the OSS Packages and your rights and responsibilities under these licenses.

In addition, the software described in this manual includes open source software developed by the Apache Software Foundation, The Legion of the Bouncy Castle, Free Software Foundation, Inc., and OpenPegasus.

Other included software is under license agreements with Hewlett-Packard Development Company, L.P.; IBM Corp.; EMC Corporation; Symantec Corporation; and The Open Group.

In addition, the software described in this manual includes open source software developed by: Copyright (c) 2005-2008, Kirill Grouchnikov and contributors All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Kirill Grouchnikov and contributors nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

C SANiQ TCP and UDP Port Usage

A P4000 SAN uses a list of well-known TCP/UDP ports to operate, see [Table 70 \(page 248\)](#).

Management applications in SAN/iQ software

Management applications include the Centralized Management Console and the scripting interface. These applications all use the ports described as management applications in the description column.

Networking best practices With SAN/iQ software

For information on networking best practices, see *Application Note - Building High Performance High Availability IP Storage Networks with SANiQ* at the following link on the Resource Center:
<http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01750150/c01750150.pdf>

Table 70 TCP/UDP ports used for normal SAN operations with SAN/iQ

IP Protocol	Port(s)	Name	Description
TCP	22	SSH	Secure Shell access for SAN/iQ Support only. Not required for normal day-to-day operations.
TCP	25	SMTP	Outgoing from storage systems for email based alerts.
TCP, UDP	53	DNS	Outgoing from storage systems for hostname resolution.
TCP, UDP	67, 68	bootps	Used only if DHCP is enabled.
UDP	123	NTP	Outgoing from storage systems for time synchronization.
UDP	161, 162	SNMP	SNMP Agent on the storage system. Incoming from SNMP clients. Outgoing from the storage systems for SNMP Traps.
UDP	514	Syslog	Outgoing from the storage systems for remote logging.
TCP, UDP	3205	iSNS	Outgoing from the storage systems for iSNS services
TCP	3260	iSCSI	iSCSI initiators connect on this port.
UDP	8453, 8454, 8455	SAN/iQ Internal	Used for internal SAN/iQ discovery.
TCP	11120	SAN/iQ Upgrade	Incoming to storage systems. Used for upgrading the storage systems. Not required for normal day-to-day operations.
TCP, UDP	13838, 13841, 13843	SAN/iQ Internal	Outgoing from management applications. Incoming to

Table 70 TCP/UDP ports used for normal SAN operations with SAN/iQ (continued)

IP Protocol	Port(s)	Name	Description
			storage systems. Used for management and control.
TCP	13840, 13850, 13851	SAN/iQ Internal	Outgoing from management applications. Incoming to storage systems. Used for management and control.
TCP	13846, 13850, 13946, 13847	SAN/iQ Support	Trace Ports for SAN/iQ Support only. Not required for normal day-to-day operations.
TCP	13847	SAN/iQ Internal	Used for Virtual Manager communication
TCP	13848	SAN/iQ Internal	Used for internal data distribution and resynchronization.
TCP	13849	iSCSI	iSCSI initiators connect to this port when using the P4000 DSM for MPIO.
UDP	27491	Console Discovery	Outgoing from management applications. Incoming to storage systems. Used by management applications to discover storage systems.
TCP, HTTP/HTTPS	2003, 5988, 5989	CIM Server	Used for HTTP requests to the CIM gateway. Outgoing from management applications. Incoming to storage systems. Used for management control. Ports used for Performance Dashboard, Global Config request to Mgmt Gateway.
TCP	13990 - 13999	CIM	Ports for CMC notifications for Alarms Events from Mgmt Gateway
UDP	13891, 13893	SAN/iQ Database	Used for heart-beating and distributing configuration changes between management nodes.
UDP	13888, 13889	SAN/iQ Control	Used for internal control communication.
UDP	14000 – 140xx	SAN/iQ Internal	Used as iSCSI targets where xx is the number of initiators
TCP	13887, 13892	Failover Manager	Communication to and from the Failover Manager, when applicable.
UDP	13947	Failover Manager	Communication to and from the Failover Manager, when applicable.
TCP	4443	SAN/iQ Diagnostics	Used by the SAN/iQ Diagnostic tool -DiagiQ that is run by P4000 Support to get advanced diagnostics

Table 70 TCP/UDP ports used for normal SAN operations with SAN/iQ *(continued)*

IP Protocol	Port(s)	Name	Description
			information using a web interface.
TCP	2301, 2302, 2381, 2382	HP System Insight Manager	Used by the HP System Insight Manager tool to display configuration information of HP platforms like DL380, DL320S using a web interface.
SSH	16022		Ports used for connecting to the cliq
TCP	11120	SAN/iQ Upgrade	Incoming to storage systems. Used for upgrading the storage systems. Not required for normal day-to-day operations.
TCP, UDP	13838, 13841, 13843	SAN/iQ Internal	Outgoing from management applications. Incoming to storage systems. Used for management and control.
TCP	13840, 13850, 13851	SAN/iQ Internal	Outgoing from management applications. Incoming to storage systems. Used for management and control.
UDP	27491	Console Discovery	Outgoing from management applications. Incoming to storage systems. Used by management applications for node discovery.

Glossary

The following glossary provides definitions of terms used in the SAN/iQ software and the HP P4000 SAN Solution.

acting primary volume	The remote volume, when it assumes the role of the primary volume in a failover scenario.
Active-Passive	A type of network bonding which, in the event of a NIC failure, causes the logical interface to use another NIC in the bond until the preferred NIC resumes operation. At that point, data transfer resumes on the preferred NIC.
Adaptive Load Balancing	A type of network bonding in which the logical interface performs load balancing of data transmission.
application-managed snapshot	Snapshot of a volume that is taken while the application that is serving that volume is quiesced. Because the application is quiesced, the data in the snapshot is consistent with the application's view of the data. That is, no data was in flight or cached waiting to be written.
authentication group	For release 7.0 and earlier, identifies the client or entity accessing the volume. Not used in release 8.0 and later.
Auto Discover	A feature in the CMC that automatically searches for storage systems on the subnet the CMC is connected to. Any storage systems it discovers appear in the navigation window on the left side of the CMC.
Bond0	Interface created for network interface failover and only appears after configuring for failover.
bonding	Combining physical network interfaces into a single logical interface.
boot device	Compact flash cards from which the storage system boots up. Also known as disk-on-modules or DOMs.
CHAP	Challenge-Handshake authentication protocol.
CLI	Command-line interface. An interface comprised of various commands which are used to control operating system responses.
clone point	The snapshot that has two or more volumes related to it. A clone point is created when a SmartClone volume is created from a snapshot or from snapshot temporary space.
cluster	A cluster is a grouping of storage systems that create the storage pool from which you create volumes.
CMC	Centralized Management Console. See HP P4000 Centralized Management Console.
communication mode	The unicast communication among storage systems and application servers.
community string	The community string acts as an authentication password. It identifies hosts that are allowed read-only access to the SNMP data.
Configuration Summary	The Configuration Summary displays an overview of the volumes, snapshots, storage systems, and iSCSI sessions in the HP StorageWorks P4000 SAN Solution. It provides an overview of the storage network broken out by management groups.
data center	Also known as a "Site." A data center is a physical location in your environment where application servers, SAN storage and network equipment reside. In the SAN/iQ Multi-Site software, a data center is typically referred to as a site.
disaster recovery site	Similar to a secondary site, the disaster recovery site is used to operate the SAN in the event of a disaster.
disk status	Whether the disk is: <ul style="list-style-type: none">• Active - on and participating in RAID• Uninitialized or Inactive - On but not participating in RAID• Off or Missing - Not on• DMA Off - disk unavailable due to faulty hardware or improperly seated in the chassis

DSM	Device Specific Module.
DSM for MPIO	The HP P4000 DSM for MPIO vendor-specific DSM that interfaces with the Microsoft MPIO framework.
failback	After failover, the process by which you restore the primary volume and turn the acting primary back into a remote volume.
failover	The process by which the user transfers operation of the application server over to the remote volume. This can be a manual operation, a scripted operation, or VMware enabled.
Failover Manager	A specialized manager running as a VMware appliance that allows you to place a quorum tie-breaker system into a 3rd location in the network to provide for automated failover/failback of the Multi-Site SAN clusters. The Failover Manager is designed to run on VMware ESX Server, VMware Server, and VMware Player. It is installed on hardware separate from the SAN hardware.
failover recovery	After failover, the process by which the user chooses to fail back to the primary volume or to make the acting primary into a permanent primary volume.
frame size	The frame size specifies the size of data packets that are transferred over the network.
full provisioning	Full provisioning reserves the same amount of space on the SAN as is presented to application servers.
ghost storage system	When using Repair Storage System, a "ghost" storage system acts as a placeholder in the cluster, keeping the cluster intact, while you repair or replace the storage system.
Graphical Legend	Describes all the icons used in the CMC: <ul style="list-style-type: none"> • Items tab - displays the icons used to represent virtual items displayed in the CMC • Hardware tab - displays the icons that represent the physical storage units.
hardware reports	Hardware reports display point-in-time statistics about the performance and health of the storage system, its drives, and configuration.
hostname	The hostname on a storage system is the user-definable name that displays below the storage system icon in the network window. It is also visible when the users browse the network.
HP StorageWorks P4000 Centralized Management Console	Management interface for the SAN/iQ software.
ID LED	LED lights on the physical storage system so that you can find that system in a rack.
iSCSI	Internet small computer system interface. Like an ordinary SCSI interface, iSCSI is standards-based and efficiently transmits block-level data between a host computer (such as a server that hosts Exchange or SQL Server) and a target device (such as the HP All-in-One Storage System). By carrying SCSI commands over IP networks, iSCSI is used to facilitate data transfers over intranets and to manage storage over long distances.
iSCSI load balancing	Improves iSCSI performance and scalability by distributing iSCSI sessions for different volumes evenly across storage systems in a cluster.
license key	A WWN-encoded sequence that is obtained from the license key fulfillment website.
Link Aggregation Dynamic Mode	A type of network bonding in which the logical interface uses both NICs simultaneously for data transfer.
log files	Log files for the storage system are stored both locally on the storage system and are also written to a remote log server.
logical site	This site is on an isolated network and power connection than the other sites. However, it can be in the same physical location as one of the real sites. Also, a site for a Failover Manager.
management group	A collection of one or more storage systems which serves as the container within which you cluster storage systems and create volumes for storage.
managers	Manager software runs on storage systems within a management group. You start managers on designated storage systems to govern the activity of all of the storage systems in the group.
MIB	Management information base. A database of managed objects accessed by network management protocols. An SNMP MIB is a set of parameters that an SNMP management station can query or set in the SNMP agent of a network device (for example, a router).

Multi-Site cluster	<p>A cluster of storage that spans multiple sites (up to three). A Multi-Site cluster must meet at least one of the following conditions:</p> <ul style="list-style-type: none"> • Contain storage systems that reside in two or more sites • Contain storage systems that span subnets • Contain multiple VIPs. The cluster can have a single site, and the multiple VIPs make it a multi-site cluster.
network RAID	Synchronous replication, mirroring or parity protection on a volume-by-volume basis. Protecting data for a volume across all storage systems in the cluster. Network RAID-10, 10+1 or 10+2 is required to protect data in an HP P4000 SAN solution.
network window	Graphically depicts the status of each storage system. Storage systems on the network are either available or part of a management group.
NTP	Network Time Protocol
original primary volume	The primary volume that fails and then is returned to service.
overprovisioned cluster	An overprovisioned cluster occurs when the total provisioned space of all volumes and snapshots is greater than the physical space available on the cluster. This can occur when there are snapshot schedules and/or thinly provisioned volumes related to the cluster.
parity	In RAID 5, redundant information is stored as parity distributed across the disks. Parity allows the storage system to use more disk capacity for data storage.
peer site	Absence of a primary site designation makes all the sites peer sites.
point-in-time snapshot	Snapshots that are taken at a specific point in time, but an application writing to that volume may not be quiesced. Thus, data may be in flight or cached and the actual data on the volume may not be consistent with the application's view of the data.
preferred interface	A preferred interface is the interface within an active backup bond that is used for data transfer during normal operation.
primary site	A site designation assigned by the administrator in the HP P4000 Centralized Management Console. A primary site is more important than a secondary site. In this setup, you would run a majority of managers in the primary site. In a two-site setup, this allows the primary site to stay online even if the network link between the primary and secondary sites fails. Typically, the primary site has majority/all of the application servers. In configurations that do not designate a primary site, the sites are referred to as "peer" sites.
primary snapshot	A snapshot of the primary volume which is created in the process of creating a remote snapshot. The primary snapshot is located on the same cluster as the primary volume.
primary volume	The volume which is being accessed (read/write) by the application server. The primary volume is the volume that is backed up with Remote Copy.
quorum	A majority of managers required to be running and communicating with each other in order for the SAN/iQ software to function.
RAID device	RAID (originally redundant array of inexpensive disks, now redundant array of independent disks) refers to a data storage scheme using multiple hard drives to share or replicate data among the drives.
RAID levels	<p>Type of RAID configuration:</p> <ul style="list-style-type: none"> • RAID 0 - data striped across disk set • RAID 1 - data mirrored from one disk onto a second disk • RAID 10 - mirrored sets of RAID 1 disks • RAID 5 - data blocks are distributed across all disks in a RAID set. Redundant information is stored as parity distributed across the disks. • RAID 50 - mirrored sets of RAID 5 disks.
RAID quorum	Number of intact disks required to maintain data integrity in a RAID set.
RAID rebuild rate	The rate at which the RAID configuration rebuilds if a disk is replaced.

RAID status	<p>Condition of RAID on the storage system:</p> <ul style="list-style-type: none"> • Normal - RAID is synchronized and running. No action is required. • Rebuild - A new disk has been inserted in a drive bay and RAID is currently rebuilding. No action is required. • Degraded - RAID is not functioning properly. Either a disk needs to be replaced or a replacement disk has been inserted in a drive. • Off - Data cannot be stored on the storage system. The storage system is offline and flashes red in the network window.
register	Register individual storage systems to use add-on applications. Registration requires sending in the storage system serial numbers to purchase the license keys, which are then applied to the storage system.
remote copy pair	The primary volume and its related remote volume.
remote snapshot	An identical copy of a primary snapshot. The remote snapshot is located on the same cluster as the remote volume.
remote volume	<p>The volume that resides in the Remote Copy location where the remote snapshots are created. The remote volume contains no data. It acts as a pointer to tell the system where to make the copy of the primary snapshot. The remote volume can be stored in these ways:</p> <ul style="list-style-type: none"> • In the same cluster in the same management group • In a different cluster in a different management group • In a different cluster in the same management group
Repair storage system	Creates a placeholder in the cluster, in the form of a “ghost” storage system, that keeps the cluster intact while you remove the storage system to replace a disk or replace the storage system itself, and return it to the cluster.
replication level	In Release 8.5 this changes to data protection level. Prior to release 8.5, replication level is the term that designated how many copies of data to keep in the cluster.
replication priority	Removed in Release 8.5. Prior to Release 8.5, replication priority allowed you to designate whether data availability or redundancy is more important in your configuration. Release 8.5 forward defaults to availability. This default can be changed using the Cliq Command Line Interface.
restripe	Striped data is stored across all disks in the cluster. You might change the configuration of a volume, for example, change data protection level, add a storage system, or remove a storage system. Because of your change, the pages in the volume must be reorganized across the new configuration. The system can keep track of several configuration changes at once. This means you can change configurations, even while a volume is in the midst of a different reconfiguration. In particular, if a reconfiguration was done by accident, you don't have to wait until it finishes to change back to the original configuration. See “Stripe”.
resync	When a storage system goes down, and writes continue to a second storage system, and the original store comes back up, the original storage system needs to recoup the exact data captured by the second storage system.
rolling back	Replaces the original volume with a read/write copy of a selected snapshot. Starting with release 8.0, the new volume retains the same name.
SAN/iQ interface	When you initially set up a storage system using the Configuration Interface, the first interface that you configure becomes the interface used for the SAN/iQ software communication.
secondary site	A site that is less important than the primary site. In this setup a minority of managers runs in the secondary site. In a two-site setup, this allows the secondary site to go offline if the network link between the Primary and secondary sites fails. Typically, the secondary site has a minority, or none, of the application servers. If the primary site fails, customers can manually recover quorum in the secondary site.
server	An application server that you set up in a management group and then assign volumes to it to provide access to those volumes.

shared snapshot	Shared snapshots occur when a clone point is created from a newer snapshot that has older snapshots below it in the tree. All the volumes created from the clone point will display these older snapshots that they share, as well as the clone point.
site	A user-designated location in which storage systems are installed. Multi-Site SAN configurations have multiple sites with storage systems in each site, and each site has its own subnet. A site can be a logical configuration, such as a subnet within the same data center, department, or application.
SmartClone volume	SmartClone volumes are space-efficient copies of existing volumes or snapshots. They appear as multiple volumes that share a common snapshot, called a clone point. They share this snapshot data on the SAN.
snapshot set	Application-managed snapshots created for a volume set.
snapshot	A fixed version of a volume for use with backup and other applications.
SNMP traps	Use traps to have an SNMP tool send alerts when a monitoring threshold is reached.
solution pack	HP P4000 Windows Solution Pack
split mirror	A full copy of data that has been split off from the original and is no longer being updated.
standard cluster	Also known as a "cluster." A standard cluster is one that does not use any of the Multi-Site features within the SAN/iQ software. Standard clusters: <ul style="list-style-type: none"> • Cannot contain storage systems that are designated to reside in a site. • Cannot contain storage systems that span subnets. • Can only have a single VIP.
storage server	Storage server software maintains the customer's data. It reads to and writes from disks in response to customer reads and writes of SANiQ volumes.
stripe	Striped data is stored across all disks in the array, which increases performance but does not provide fault tolerance.
synchronize	The process of copying the most recent snapshot from the primary volume to a new remote snapshot. On failback, synchronization is the process of copying the most recent remote snapshot back to the primary volume. The CMC displays the progress of this synchronization. Also, you can manually synchronize if necessary to include data that is on the remote volume but not the primary.
target secret	Target secret is used in both 1-way and 2-way CHAP when the target (volume) challenges the iSCSI initiator.
temporary space	Temporary space is created when a snapshot is mounted for use by applications and operating systems that need to write to the snapshot when they access it. Temporary space can be converted to a volume using the SmartClone process.
thin provisioning	Thin provisioning reserves less space on the SAN than is presented to application servers.
trap	A type of SNMP message used to signal that an event has occurred. (SNIA)
Trap Community String	The Trap Community String is used for client-side authentication when using SNMP.
unicast	Communication between a single sender and a single receiver over a network.
VIP	virtual IP address
virtual IP address	A highly available address that ensures that if a storage system in a cluster becomes unavailable, servers can still access the volume through the other storage systems in the cluster.
virtual machine	A virtual storage appliance that provides one or more simultaneous storage environments in which SAN/iQ may execute as though they were running on the bare iron.
virtual manager	A manager that is added to a management group but is not started on a storage system until it is needed to regain quorum.
volume	A logical entity that is made up of storage on one or more storage systems. It can be used as raw data storage or it can be formatted with a file system and used by a host or file server.
volume lists	For release 7.0 and earlier, provide the link between designated volumes and the authentication groups that can access those volumes. Not used in release 8.0 and later.

volume set	Two or more volumes used by an application. For example, you may set up Exchange to use two volumes to support a StorageGroup: one for mailbox data and one for logs. Those two volumes make a volume set.
volume size	The size of the virtual device communicated to the operating system and the applications.
VSS Provider	HP P4000 VSS Provider is the hardware provider that supports the Volume Shadow Copy Service on the HP P4000 SAN Solution.
VSS	Volume Shadow Copy Service
writable space	See temporary space .

Index

Symbols

1000BASE T interface, 53

A

access control

SNMP, 93

access rights see permission levels

accessing

volumes from servers, 196

active interface

in active-passive bond, 57

in adaptive load balancing bond, 62

in link aggregation dynamic mode bond, 60

Active-Passive bond, 57

active interface, 57

during failover, 58

example configurations, 58

requirements, 57

Adaptive Load Balancing bond

active interface, 62

during failover, 63

example configurations, 63

preferred interface, 62

requirements, 62

add-on applications

overview, 222

registering for, 224

adding

a remote log, 101

administrative groups, 80

administrative users, 79

clusters, 132

DNS domain name, 71

DNS servers, 71

domain names to DNS suffixes, 72

iSNS server, 133

management groups, 104

managers to management group, 114

requirements for snapshots, 162

requirements for volumes, 154

routes, 73

servers to management groups, 197

snapshots, 163

SNMP clients, 93

statistics, 216

storage for the first time, 16

storage systems to existing cluster, 135

storage systems to management group, 114

storage to clusters, 135

users to a group, 82

virtual manager, 129

volumes, 156

administrative groups, 80

adding, 80

adding users, 82

changing, 81

deleting, 82

permission levels, 81

permissions descriptions, 81, 82

removing users, 82

administrative security, 103

administrative users, 79

adding, 79

deleting, 80

agents

disabling SNMP, 94

enabling for SNMP, 92

alarms

displaying details, 87

displaying in a separate window, 87

exporting, 88

filtering, 87

overview, 85

window for viewing, 13

working with, 87

analyzer

Best Practice, 111

Configuration, 108

application servers, clustering, 196, 199

application-managed snapshots

converting temporary space from, 172

creating, 164, 167

creating for volume sets, 164

creating schedules for volume sets, 166

creating SmartClone volumes from, 175

defined, 161

deleting, 176

making available, 171

requirements for, 163

rolling back from, 174, 175

Assign Volume and Snapshot wizard

servers, 18

assigning servers to volumes and snapshots, 202, 203, 204

authentication groups

and volume lists, 196

auto discover, 13

systems on network, 19

turning off, 19

auto performance protection, 136

storage system inoperable, 137

storage system overloaded, 137

volume availability and, 137

availability of volumes and snapshots, 28, 137

Availability tab, 28

available storage systems see available systems

available system pool, 103

available systems

defined, 14

upgrading software, 26

B

- backing out
 - of Remote Copy evaluation, 223
 - of scripting evaluation, 224
- bandwidth, changing local settings, 116
- best practice
 - configuring cluster for disaster recovery, 127
 - frame size, 52
 - link aggregation dynamic mode, 56
 - network, 48
 - NIC bonds, 64
 - recommended numbers for management group storage items, 109
 - setting volume size, 141
 - speed and duplex settings, 51
 - using snapshots for protection against data deletion, 161
- Best Practice Summary
 - cluster-level data protection, 112
 - disk level data protection, 112
 - disk protection using RAID, 112
 - disk RAID consistency, 112
 - large single-system SATA cluster, 112
 - network bond consistency, 113
 - network bonding, 113
 - network flow control consistency, 113
 - network frame size consistency, 113
 - overview, 111
 - systems running managers, 113
 - volume access, 113
 - volume-level data protection, 113
- block device, iSCSI as, 151
- block storage system, 151
- boot devices
 - checking status of dedicated, 29
 - dedicated, 29
 - status of dedicated, 29
- Boot Devices tab, 29
- BOOTP, 55

C

- capacity
 - clusters, 132
 - clusters and usable space in, 150
 - disk capacity and volume size, 151
 - of the SAN, 140
 - planning thin provisioning, 141
 - planning volume size, 140
 - planning, full provisioning, 141
 - storage systems, 132
- capacity management
 - and deleting snapshots, 147
 - and scheduled snapshots, 147, 165
 - monitoring SAN for, 147
 - viewing SAN capacity and usage, 147
 - volume size and snapshots, 147
- caution
 - logging into management group, 13
- Centralized Management Console

- features of, 12
- finding systems on network, 19
- graphical legend, 13
- icons used in, 13
- menu bar, 13
- navigation window, 19
- overview, 12
- setting up for remote support, 20
- turning off auto discover in, 19
- upgrading, 25, 26

Challenge Handshake Authentication Protocol *see* CHAP

- changing
 - administrative group description, 81
 - administrative group permissions, 81
 - changing RAID erases data, 36
 - cluster configuration, 133
 - clusters for volumes, 159
 - data protection levels, 159
 - host names, 22
 - IP address of storage system, 54
 - local bandwidth, 115
 - maintenance mode to normal, 118
 - management groups, 115
 - network configuration, 49
 - order of NTP server access, 77
 - snapshots, 165
 - thresholds in a snapshot, 165
 - user password, 79
 - volume descriptions, 158
 - volume size, 159
- CHAP
 - 1-way, 230
 - 2-way, 230
 - editing, 199
 - iSCSI, 230
 - requirements for configuring, 198, 231
 - terminology in different initiators, 231
 - using, 230
 - volumes and, 230
- characteristics of SmartClone volumes, 180
- checklist for disk replacement, 44
- choosing a RAID configuration, 31
- clearing
 - items in navigation window, 19
 - statistics sample data, 218
- client access to volumes using Assign Volume and Snapshot wizard, 18
- clients, adding SNMP, 93
- clone *see* SmartClone volumes
- clone a volume, 179
- clone point
 - and shared snapshots, 186
 - deleting, 193
 - SmartClone volumes, 185
- cluster-level data protection
 - Best Practice Summary, 112
- clustered application servers
 - changing volume associations after deleting server cluster, 202

- creating, [200](#)
 - deleting cluster, [201](#)
 - editing, [201](#)
- clustering application servers, [196](#), [199](#)
 - requirements for, [199](#)
- clustering managers, [103](#)
- clustering servers, [196](#), [199](#)
- clusters
 - adding, [132](#)
 - adding storage system, [135](#)
 - capacity, [132](#)
 - changing volumes, [159](#)
 - comparing the load of two, [209](#), [215](#)
 - data protection levels, [142](#)
 - deleting, [139](#)
 - editing, [133](#)
 - map view, [133](#)
 - overview, [132](#)
 - removing storage systems from, [136](#)
 - reorder storage systems within, [136](#)
 - repairing storage system in, [138](#)
 - space reporting in, [148](#)
 - swapping storage systems, [135](#)
 - System Use window, [150](#)
 - troubleshooting, [136](#)
 - usable space in, [150](#)
 - Use Summary window, [148](#)
 - viewing capacity and usage of, [147](#)
 - Volume Use window, [149](#)
- CMC *see* Centralized Management Console
- communication interface for SAN/iQ communication, [74](#)
- compliant iSCSI initiators, [230](#)
- configuration
 - best practice summary, [111](#)
 - changing network, [49](#)
- configuration categories
 - storage system, defined, [21](#)
 - storage systems, [21](#)
- Configuration Interface
 - configuring frame size in, [237](#)
 - configuring network connection in, [236](#)
 - configuring TCP speed and duplex in, [237](#)
 - connecting to, [235](#)
 - creating administrative users in, [236](#)
 - deleting NIC bond in, [237](#)
 - resetting DSM configuration in, [238](#)
 - resetting storage system to factory defaults in, [238](#)
- Configuration Summary
 - configuration guidance, [109](#)
 - management group summary roll-up, [108](#)
 - overview, [108](#)
 - reading for management group, [110](#)
- configurations
 - RAID, [31](#)
- configuring
 - disabled network interface, [71](#)
 - frame size in Configuration Interface, [237](#)
 - IP address manually, [54](#)
 - iSCSI single host, [232](#)
 - network connection in Configuration Interface, [236](#)
 - network interface bonds, [64](#)
 - network interfaces, [54](#)
 - NIC speed and duplex, [50](#)
 - RAID, [31](#)
 - split network, [48](#)
 - storage systems, [17](#)
 - TCP speed and duplex in Configuration Interface, [237](#)
 - virtual IP address, [132](#)
 - virtual manager, [129](#)
- connecting to the Configuration Interface, [235](#)
- consumed space by volume, [149](#)
- contacting HP, [239](#)
- converting temporary space
 - from application-managed snapshots, [172](#)
- copying, volumes from snapshots, [169](#)
- creating *see* adding
 - administrative users in Configuration Interface, [236](#)
 - new administrative group, [82](#)
 - NIC bond, [65](#)
 - server cluster, [200](#)
 - SmartClone volumes, [188](#)
 - storage, [140](#)
 - volumes using the wizard, [17](#)
- critical events
 - defined, [85](#)
- CSV file, exporting performance statistics to, [220](#)
- custom event filters
 - creating, [89](#)
 - deleting, [90](#)
- customer self repair, [240](#)
- customer support
 - registering Remote Copy, [224](#)
- D**
- data
 - and deleting volumes, [159](#)
 - clearing statistics sample, [218](#)
 - stripe patterns in clusters, [136](#)
- data mining using SmartClone volumes, [179](#)
- data protection
 - changing levels for volumes, [159](#)
 - requirements for setting levels, [155](#)
- data protection level
 - allowed in clusters, [142](#)
 - for volumes, [142](#)
 - Network RAID-10, [143](#)
 - Network RAID-10+2, [144](#)
 - Network RAID-6, [145](#)
 - Network RAID-10+1, [143](#)
 - repairing storage system in cluster with data protection level, [138](#)
 - requirements for setting, [155](#)
- data reads and writes and RAID status, [36](#)
- data redundancy
 - and RAID status, [36](#)
- data transmission, [52](#)
- date
 - setting with NTP, [76](#)

- setting without NTP, 78
- date and time for scheduled snapshot, 167
- decreasing volume size, 159
- defaults
 - font size in CMC, 15
 - language displayed in CMC, 15
 - naming conventions in CMC, 15
 - restoring for the Performance Monitor, 218
- definition
 - RAID configurations, 31
 - SmartClone volumes, 177
- degraded RAID, 36
- deleting
 - administrative groups, 82
 - administrative users, 80
 - an administrative group, 82
 - clone point, 193
 - clusters, 139
 - custom event filters, 90
 - DNS servers, 72
 - management groups, 119
 - multiple SmartClone volumes, 194
 - network interface bonds, 68
 - Network RAID-5 and Network RAID-6 snapshots, space considerations for, 176
 - NIC bond in Configuration Interface, 237
 - NTP server, 77
 - prerequisites for volumes, 173, 176
 - restrictions on for snapshots, 176
 - restrictions on for volumes, 159
 - routing information, 73
 - server cluster, 201
 - server cluster and change volume associations, 202
 - servers, 199
 - SmartClone volumes, 193
 - snapshot schedules, 168
 - snapshots, 175
 - snapshots, and capacity management, 147
 - volumes, 159
- descriptions
 - changing for clusters, 133
 - changing for volumes, 158
- Details tab
 - storage systems, 25
- details, viewing for statistics, 217
- DHCP
 - using, 55
 - warnings when using, 55
- diagnostics
 - hardware, 96
 - list of diagnostic tests, 97
 - viewing reports, 96
- disabled network interface, configuring, 71
- disabling
 - network interfaces, 70
 - SNMP agent, 94
 - SNMP traps, 95
- disassociating management groups, 115
 - see also *HP P4000 Remote Copy User Guide*
- disaster recovery
 - best practice, 127
 - starting virtual manager, 130
 - using a virtual manager, 125
- disk
 - arrangement in storage systems, 39
 - disk setup report, 37
 - managing, 37
 - managing in storage system, 37
 - powering off through the CMC, 45
 - powering on through the CMC, 45
 - replacement, 45, 46
 - replacement checklist, 44
 - replacing in RAID 1 or RAID 10, and RAID 5, and RAID 50, 45
 - replacing in replicated cluster, 138
 - replacing in storage system, 42
 - using Repair Storage System when replacing, 43
 - VSA, recreating, 40
- disk capacity in cluster, 150
- disk drive see disk
- disk level data protection
 - Best Practice Summary, 112
- disk protection using RAID
 - Best Practice Summary, 112
- disk RAID consistency
 - Best Practice Summary, 112
- disk report, 38
- disk setup
 - DL320s, 39
 - P4300, 41
 - P4500, 40
 - report, 38
 - tab, 38
- disk space usage, 151
- disk status
 - DL320s, 39
 - P4500, 40
 - P4800, 42
 - VSA, 40
- display
 - default naming conventions, 15
 - font size, 15
 - languages, 15
- display tools
 - for using map view, 14
- DL320s
 - disk setup, 39
 - disk status, 39
- DNS
 - adding domain name, 71
 - and DHCP, 71
 - and static IP addresses, 71
- DNS server
 - adding, 71
 - and static IP addresses, 71
 - editing IP or domain name, 72
 - removing, 72
 - using, 71

- document
 - related information, [240](#)
- documentation
 - HP website, [240](#)
- Domain Name Server *see* DNS Server
- domain names
 - adding to DNS suffixes, [72](#)
 - editing in DNS suffixes list, [72](#)
 - removing from DNS suffixes list, [72](#)
- downloading
 - upgrades, [25](#)
- DSM for MPIO, [229](#)
 - how to see if using, [234](#)
 - tips for using to access volumes from servers, [234](#)
 - when using two NICs, [204](#)
- duplex, configuring, [50](#)
- Dynamic Host Configuration Protocol *see* DHCP

E

- editing
 - clusters, [133](#)
 - DNS server domain names, [72](#)
 - DNS server IP addresses, [72](#)
 - domain name in DNS suffixes list, [72](#)
 - frame size, [52](#)
 - group name, [81](#)
 - management groups, [115](#)
 - network interface frame size, [51](#)
 - network interface speed and duplex, [50](#)
 - NTP server, [77](#)
 - routes, [73](#)
 - server cluster, [201](#)
 - servers, [199](#)
 - SmartClone volumes, [193](#)
 - snapshot schedules, [167](#)
 - snapshots, [165](#)
 - SNMP trap recipient, [94](#)
 - volumes, [157](#)
- email
 - setting up for event notification, [91](#)
- enabling
 - NIC flow control, [53](#)
 - SNMP traps, [94](#)
- establishing network interfaces, [53](#)
- eth0 and eth1, [53](#)
- Ethernet interfaces, [53](#)
- evaluating
 - backing out of Remote Copy, [223](#)
 - backing out of scripting, [224](#)
 - Remote Copy, [222](#)
 - scripting, [223](#)
- event notification
 - configuring access control for SNMP clients, [93](#)
 - enabling SNMP agents for, [92](#)
 - setting up email for, [91](#)
 - setting up email recipients for, [91](#)
 - setting up email server for, [91](#)
 - setting up SNMP for, [92](#)
- events

- changing the retention period, [88](#)
- copying to the clipboard, [90](#)
- deleting custom filters, [90](#)
- displaying details, [90](#)
- displaying in a separate window, [88](#)
- displaying new, [89](#)
- exporting, [90](#)
- filtering, [89](#)
- overview, [85](#)
- setting up remote log destinations, [88](#)
- types defined, [85](#)
- working with, [88](#)
- events list
 - column descriptions, [86](#)
- example scenarios for using SmartClone volumes, [178](#)
- exchanging
 - storage systems in clusters, [136](#)
- exporting
 - alarms, [88](#)
 - events, [90](#)
 - performance data, [220](#)
 - performance statistics to a CSV file, [220](#)
 - support logs, [102](#)

F

- Failover Manager, [107](#), [108](#)
 - and Multi-Site SAN, [107](#)
 - requirements for, [120](#)
- fault tolerance, [229](#)
 - data protection level for volumes, [142](#)
 - network interface bonding, [55](#)
 - quorum and managers, [106](#)
 - stopping managers, [115](#)
- faults, isolating, [206](#)
- Feature Registration tab, [224](#), [226](#)
- features
 - of Centralized Management Console, [12](#)
- file systems, [151](#)
 - mounting on volumes, [154](#)
- filtering
 - alarms, [87](#)
 - events, [89](#)
- finding storage systems, [17](#)
 - auto discover, [13](#)
 - on the network, [19](#)
 - troubleshooting, [19](#)
- first storage system, [104](#)
- flow control, [52](#)
- font size
 - setting in preferences, [15](#)
- formatting volumes for use, [204](#)
- frame size
 - NIC, [51](#)
 - VSA, [48](#)
- frames, editing size, [52](#)
- full permissions, [81](#)

G

- gateway session for VIP with load balancing, [230](#)

- Getting Started Launch Pad, 16
- ghost storage system, 138
 - removing after data is rebuilt, 246
- Gigabit Ethernet, 53
 - see also GBe
- glossary
 - for SAN/iQ software and HP P4000 SAN, 251
 - SmartClone volumes, 177
- graphical legend
 - Centralized Management Console, 13
- group name
 - editing, 81
- groups
 - administrative, 80
 - administrative default groups, 80
 - deleting administrative, 82

H

- hardware diagnostics, 96
 - list of diagnostic tests, 97
 - tab window, 96
- hardware information report, 97
 - saving to a file, 98
- help
 - obtaining, 239
- highlighting lines, 219
- host names
 - access SNMP by, 93
 - changing, 22
 - resolution, 22
- host storage system for virtual IP address, 229
- hot swap
 - RAID degraded, 37
 - safe to remove status, 43
- HP
 - technical support, 239
- HP DSM for MPIO, 229

I

- I/O performance, 136
- icons
 - licensing, 223
 - used in Centralized Management Console, 13
- identifying network interfaces, 53
- increasing volume size, 159
- informational events
 - defined, 85
- Insight Remote Support software, 239
- installing
 - SNMP MIB, 95
- interface
 - administrative users in, 236
 - configuration, 235
 - configuring network connection in, 236
 - connecting to, 235
 - deleting NIC bond in, 237
 - resetting storage system to factory defaults, 238
- IP addresses
 - changing, iSNS server, 133

- configuring for storage system, 54
- NTP server, 77
- pinging, 54
- removing iSNS server, 133
- using DHCP/BOOTP, 55

iSCSI

- and CHAP, 230
- and fault tolerance, 229
- and iSNS servers, 229
- and virtual IP address, 229
- as block device, 151
- authentication, 230
- changing or removing virtual IP, 134
- CHAP, 230
- clusters and VIP, 229
- configuring CHAP, 198, 231
- load balancing, 230
- load balancing and compliant initiators, 230
- logging on to volumes, 204
- performance, 230
- setting up volumes as persistent targets, 204
- single host configuration, 232
- terminology in different initiators, 231
- volumes and, 230

- iSCSI initiators
 - configuring virtual IP addresses for, 132
- iSNS server
 - adding, 133
 - and iSCSI targets, 229
 - changing or removing IP address, 133

L

- layout of disks in storage systems, 39
- license keys, 224
- licensing icons, 223
- lines
 - changing the color of in the Performance Monitor, 219
 - changing the style of in the Performance Monitor, 219
 - displaying or hiding in the Performance Monitor, 219
 - highlighting, 219
- Link Aggregation Dynamic Mode bond, 60
 - active interface, 60
 - during failover, 61
 - example configurations, 61
 - preferred interface, 60
 - requirements, 60
- list of diagnostic tests, 97
- load balancing
 - compliant iSCSI initiators, 230
 - editing, 199
 - gateway session when using, 230
 - iSCSI, 230
- local bandwidth, setting, 115
- locale
 - setting in preferences, 15
- locating a storage system in a rack, 23
- log files
 - saving for technical support, 100
- log in

- to a storage system in a management group, 22
- to management group, 13, 113
- to storage systems in Available Systems pool, 13
- log out
 - of management group, 114
- logging on to volumes in iSCSI, 204
- logs
 - exporting support, 102

M

- maintenance mode
 - changing to normal mode, 118
 - management group in, 117
- management group
 - caution for logging into, 13
- management group time
 - refreshing, 76
- management groups
 - adding, 104
 - adding servers to, 197
 - adding storage systems, 104, 114
 - best practice recommendations, 109
 - configuration guidance, 109
 - Configuration Summary roll-up, 108
 - deleting, 119
 - editing, 115
 - function, 103
 - functions of managers, 106
 - logging in, 113
 - logging out, 114
 - maintenance mode, 117
 - normal mode, 118
 - overview, 103
 - prerequisites for removing storage systems, 118
 - reading Configuration Summary, 110
 - removing storage systems, 118
 - requirements for adding, 103
 - saving configuration information, 116
 - setting local bandwidth, 115
 - shutdown procedure, 117
 - shutting down, 116
 - starting managers, 114
 - starting up, 117
 - using virtual manager configuration for, 125
 - using virtual manager in disaster recovery, 125
- Management Information Base see MIB
- manager IP addresses
 - updating, 75
- managers
 - Failover, 107, 108
 - functions of, 106
 - implications of stopping, 115
 - overview, 106
 - quorum and fault tolerance, 106
 - starting, 114
 - stopping, 115
 - virtual, 125
- managers on storage systems
 - Best Practice Summary, 113

- managing disks, 37
- Map View, 189
 - for SmartClone volumes, 190
- map view
 - changing the view, 14
 - display tools, 14
 - for clusters, 133
 - for volumes, 157
 - possible views and layouts for network elements, 15
 - using, 14
- menu bar
 - Centralized Management Console, 13
- MIB
 - for SNMP, 95
 - installing, 95
 - locating, 95
 - versions, 95
- migrating RAID, 35
- mixed RAID, 35
- monitoring
 - performance, 205
 - RAID status, 36
- monitoring interval in the Performance Monitor, 216
- Motherboard Port1 and Motherboard Port2, 53
- mounting snapshots, 169
- move
 - storage systems within cluster, 136
- Multi-Site SAN
 - and Failover Manager, 107

N

- naming conventions
 - setting in preferences, 15
- naming SmartClone volumes, 180, 181
- navigation window
 - clearing items in, 19
 - overview, 13
- network
 - best practices, 48
 - change configuration of, 49
 - finding storage systems on, 13, 19
 - managing settings, 48
 - overview, 48
 - ping IP address, 54
 - split configurations for, 48
- network bond consistency
 - Best Practice Summary, 113
- network bonding
 - Best Practice Summary, 113
- network flow control consistency
 - Best Practice Summary, 113
- network frame size consistency
 - Best Practice Summary, 113
- network interface bonds, 55
 - active-passive, 57
 - adaptive load balancing, 62
 - best practices, 64
 - communication after deleting, 69
 - configuring, 64

- creating, 65
- deleting, 68, 237
- determining if use would improve performance, 208
- link aggregation dynamic mode, 60
- physical and logical interface, 57
- requirements, 56
- requirements for Adaptive Load Balancing, 62
- setting flow control, and, 53
- settings, 56
- status of, 67
- verifying, 66
- VSA, 48
- network interfaces, 60
 - attaching Ethernet cables, 53
 - bonding, 55
 - configuring, 54, 71
 - disabling or disconnecting, 70
 - establishing, 53
 - identifying, 53
 - speed and duplex settings, 50
 - used for SAN/iQ communication, 74
 - VSA, 48
- Network RAID-10
 - description of, 143
- Network RAID-10+1
 - description of, 143
- Network RAID-10+2
 - description of, 144
- Network RAID-6
 - description of, 145
- Network Time Protocol *see* NTP
- network window *see* navigation window
- NIC *see* network interfaces
- NIC flow control, 52
 - enabling, 53
 - requirements, 53
 - VSA, 48
- normal RAID
 - status, 36
- not found
 - storage systems, 20
- not preferred NTP server, 77
- NTP
 - selecting, 76
 - server, 76
 - server, deleting, 77
 - servers, changing list order, 77

O

- off RAID
 - status, 36
- ordering NTP access list, 77
- overview
 - add-on applications, 222
 - Centralized Management Console, 12
 - clusters, 132
 - disk replacement in special cases, 241
 - management groups, 103
 - managers, 106

- network, 48
- provisioning storage, 140
- setting date and time, 76
- SmartClone volumes, 177
- snapshots, 146, 161
- SNMP, 92
- storage category, 30
- volumes, 154

P

- P4300
 - disk setup, 41
- P4500
 - disk setup, 40
 - disk status, 40
- P4800
 - disk setup, 42
 - powering off the system controller and disk enclosure, correct order, 23
 - powering on the system controller and disk enclosure, correct order, 23
- passwords
 - changing in Configuration Interface, 236
- pausing
 - monitoring, 219
- performance *see* I/O performance
- performance and iSCSI, 230
- Performance Monitor
 - current SAN activity example, 206
 - exporting data from, 220
 - fault isolation example, 206
 - learning about applications on the SAN, 207
 - learning about SAN performance, 205
 - load comparison of two clusters example, 209
 - load comparison of two volumes example, 209
 - NIC bonding example, 208
 - overview, 205
 - pausing, 219
 - planning for SAN improvements, 208
 - prerequisites, 205
 - restarting, 219
 - statistics, defined, 213
 - understanding and using, 205
 - workload characterization example, 206
- Performance Monitor graph
 - changing, 219
 - changing line color, 219
 - changing line style, 219
 - changing the scaling factor for, 220
 - displaying a line, 219
 - hiding, 219
 - hiding a line, 219
 - showing, 219
- Performance Monitor window
 - accessing, 210
 - graph, 211
 - parts defined, 210
 - saving to an image file, 221
 - table, 212

- toolbar, 211
- permissions
 - administrative group, 81
 - effect of levels, 203
 - full, 81
 - read modify, 81
 - read only, 81
- ping IP address, 54
- planning
 - RAID configuration, 33
 - SmartClone volumes, 179
 - snapshots, 146, 162
 - volume size, 140
 - volumes, 140, 154
- planning capacity
 - full provisioning method, 141
 - thin provisioning method, 141
- point-in-time snapshots
 - defined, 161
- pool of storage, 103
- positioning
 - storage systems in cluster, 136
- powering off
 - disk, using CMC, 45
 - P4800 system controller and disk enclosure, correct order, 23
 - storage systems, 24
- powering on
 - disk, using CMC, 45
 - P4800 system controller and disk enclosure, correct order, 23
- preferences
 - setting font size, 15
 - setting for upgrades, 25
 - setting locale, 15
 - setting naming conventions, 15
- preferred interface
 - in active-passive bond, 57
 - in adaptive load balancing, 62
 - in link aggregation dynamic mode bond, 60
- preferred NTP server, 77
- prerequisites
 - deleting volumes, 159
- prerequisites for
 - adding volumes, 154
 - assigning servers to volumes, 202
 - deleting volumes, 160, 173, 176
 - Performance Monitor, 205
 - removing storage systems from management group, 118
 - servers, 197
 - SmartClone volumes, 177
- primary interface, NICs, 74
- primary type see volumes
- primary volumes, 154
- protection
 - RAID vs. volume protection, 33
- protocols, DHCP, 55
- provisionable space in cluster, 148

- provisioned space in cluster, 148
- provisioning storage, 140
 - and space allocation, 140

Q

- quorum
 - and managers, 106
 - starting virtual manager to recover, 130
 - stopping managers, 115

R

- RAID
 - and data protection, 33
 - as data protection, 33
 - changing RAID erases data, 36
 - configurations defined, 31
 - configuring, 30, 31
 - definitions, 31
 - device, 32
 - device status, 30
 - managing, 31
 - planning configuration, 33
 - procedure for reconfiguring, 36
 - rebuild rate, 35
 - rebuilding, 46
 - reconfiguring, 36
 - reconfiguring requirements, 36
 - replacing a disk, 45, 46
 - replication in a cluster, 34
 - requirements for configuring, 36
 - resyncing, 137
 - setting rebuild rate, 35
 - status, 36
 - status and data reads and writes, 36
 - status and data redundancy, 36
- RAID (virtual), devices, 33
- RAID 0
 - single disk replacement, 45
- RAID 1 or RAID 10
 - single disk replacement, 45
- RAID 5, RAID 50
 - single disk replacement, 45
- RAID 6
 - single disk replacement, 45
- RAID consistency
 - Best Practice Summary, 112
- RAID levels defined, 241
- RAID status
 - status, 36
- raw space in cluster, 150
- raw storage, 151
- read only permissions, 81
- read only volumes, 169
- read-modify permissions, 81
- rebooting
 - storage systems, 24
- rebuild data
 - when not running manager, 242
- rebuild volume data, 245

- rebuilding
 - RAID, 46
 - rate for RAID, 35
- reclaimable space in volumes, 149
- reconfiguring RAID, 36
- recreate the RAID array, 243
- redundant array of independent disks *see* RAID
- registering add-on applications, 224
- registering features
 - Feature Registration tab, 224, 226
 - for a storage system, 28
- related documentation, 240
- Remote Copy
 - backing out of evaluation, 223
 - evaluating, 222
 - registering, 224
 - remote volumes, 154
- remote destinations
 - setting up for events, 88
- remote log files, 101
 - adding, 101
 - changing remote log file target computer, 101
 - configuring target computer, 101
 - removing old logs, 101
- remote support
 - setting up CMC for, 20
- remote support software, 239
- remote volumes, 154
 - see also* HP P4000 Remote Copy User Guide
- removing
 - administrative users from a Group, 82
 - DNS server, 72
 - domain name from DNS suffixes list, 72
 - ghost storage system after the data is rebuilt, 246
 - old log files, 101
 - prerequisites for storage systems from management groups, 118
 - SNMP trap recipient, 94
 - statistics, 218
 - storage systems from cluster, 136
 - storage systems from management groups, 118
 - systems from view in CMC, 19
 - users from administrative groups, 82
 - virtual manager, 131
- reorder
 - storage systems in clusters, 136
- repair storage system, 138
 - prerequisites, 138
 - replacing a disk, 43
- replacing
 - disks, 42
 - storage systems in clusters, 136
- reports
 - diagnostic, 96
 - disk, 37
 - disk setup for RAID, 38
 - Hardware Information, 97
 - RAID setup, 32
 - saving Hardware Report to a file, 98

- storage system statistics, 97
- requirements for
 - adding management group, 103
 - adding volumes, 154
 - application-managed snapshots, 163
 - changing SmartClone volumes, 192
 - changing volumes, 158
 - configuring CHAP for iSCSI, 198, 231
 - editing snapshots, 165
 - Failover Manager, 120
 - network interface bonding, 56
 - rolling back volumes, 173
 - scheduling snapshots, 165
 - server clusters, 199
 - split network, 48
 - system for Failover Manager on ESX Server, 122
 - virtual manager, 126
- resetting
 - DSM in Configuration Interface, 238
 - storage system to factory defaults, 238
- resolving host names, 22
- restarting monitoring, 219
- restoring
 - defaults for the Performance Monitor, 218
 - volumes, 172
- restriping, volume, 46
- resyncing
 - RAID, 137
- return the repaired storage system to the cluster, 244
- rolling back a volume, 172
 - from application-managed snapshots, 174, 175
 - restrictions on, 173
- routing
 - adding network, 73
 - deleting, 73
 - editing network, 73
- routing tables
 - managing, 73

S

- safe to remove status, 43
- sample interval
 - changing for the Performance Monitor, 216
- SAN
 - capacity of, 140
 - comparing the load of two clusters, 209, 215
 - comparing the load of two volumes, 209
 - current activity performance example, 206
 - determining if NIC bonding would improve performance, 208
 - fault isolation example, 206
 - learning about applications on the SAN, 207
 - learning about SAN performance, 205
 - monitoring performance, 205
 - planning for SAN improvements, 208
 - using Performance Monitor, 205
 - workload characterization example, 206
- SAN configuration
 - best practice summary, 111

- SAN/iQ
 - upgrading, [25, 26](#)
- saved space in cluster, [148](#)
- saved space in volumes, [149](#)
- saving
 - diagnostic reports, [96](#)
 - log files for technical support, [100](#)
 - management group configuration information, [116](#)
- scaling factor
 - changing, [220](#)
- scheduled snapshots
 - requirements for, [165](#)
- scripting evaluation, [223](#)
 - backing out of, [224](#)
 - turning off, [224](#)
 - turning on, [223](#)
- searching for storage systems, [13, 19](#)
- security
 - administrative, [103](#)
 - of storage resources, [103](#)
- server access
 - SmartClone volumes, [180](#)
- server cluster, [196, 199](#)
 - change volume associations after deleting, [202](#)
 - creating, [200](#)
 - deleting, [201](#)
 - editing, [201](#)
 - requirements for, [199](#)
- servers
 - access to volumes and snapshots, [196](#)
 - access using Assign Volume and Snapshot wizard, [18](#)
 - adding DNS, [71](#)
 - adding iSNS, [133](#)
 - adding to management groups, [197](#)
 - assigning to volumes and snapshots, [202, 203](#)
 - clustering, [196, 199](#)
 - deleting, [199](#)
 - editing, [199](#)
 - editing assignments to volumes and snapshots, [203, 204](#)
 - editing IP or domain name for DNS, [72](#)
 - editing NTP, [77](#)
 - NTP, [76](#)
 - preferred, not preferred for NTP, [77](#)
 - prerequisites for, [197](#)
 - prerequisites for assigning to volumes, [202](#)
 - removing DNS, [72](#)
- service bundles
 - exporting, [102](#)
- set ID LED, [23](#)
- setting
 - IP address, [54](#)
 - local bandwidth, [115](#)
 - RAID rebuild rate, [35](#)
- setting date and time, [76](#)
 - for management group, [76](#)
 - overview, [76](#)
 - procedure, [76, 78](#)
 - refreshing for management group, [76](#)
 - setting time zone, [78](#)
 - time zone on storage system, [76, 78](#)
 - with NTP, [76](#)
 - without NTP, [78](#)
- setting up a RAID disk, [38](#)
- shared snapshots, [186](#)
- shutting down a management group, [116, 117](#)
- single disk replacement in RAID 0, [45](#)
- Single Host Configuration in iSCSI, [232](#)
- single system, large cluster using SATA
 - Best Practice Summary, [112](#)
- size
 - changing for volumes, [159](#)
 - for snapshots, [147](#)
 - planning for snapshots, [162](#)
 - planning for volumes, [140](#)
 - requirements for volumes, [155](#)
- slow I/O, [136](#)
- SmartClone volumes
 - assigning server access, [180](#)
 - characteristics of, [180](#)
 - characteristics of, shared versus individual, [182](#)
 - clone point, [185](#)
 - creating from application-managed snapshots, [175](#)
 - definition of, [177](#)
 - deleting, [193](#)
 - deleting multiple, [194](#)
 - editing, [193](#)
 - examples for using, [178](#)
 - glossary for, [177](#)
 - making application-managed snapshot available after
 - creating, [171](#)
 - overview, [177](#)
 - planning, [179](#)
 - planning naming convention, [180](#)
 - planning space requirements, [179](#)
 - requirements for changing, [192](#)
 - uses for, [179](#)
 - viewing with Map View, [190](#)
- snapshots
 - adding, [163](#)
 - application-managed, [161](#)
 - as opposed to backups, [146](#)
 - assigning to servers, [202, 203](#)
 - changing thresholds, [165](#)
 - controlling server access to, [196](#)
 - copying a volume from, [169](#)
 - creating application-managed, [164, 167](#)
 - creating application-managed for volume sets, [164](#)
 - deleting, [175](#)
 - deleting Network RAID-5 and Network RAID-6, space
 - considerations for, [176](#)
 - deleting schedules, [168](#)
 - editing, [165](#)
 - editing schedules, [167](#)
 - editing server assignments, [203, 204](#)
 - managing capacity and scheduled snapshots, [165](#)
 - mounting, [169](#)
 - overview, [146, 161](#)

- planning, 146, 162
 - planning size, 162
 - point-in-time, 161
 - read/write and deleting temporary space, 172
 - requirements for editing, 165
 - restrictions on deleting, 176
 - rolling back a volume from, 172
 - shared, 186
 - size, 147
 - temporary space for read/write snapshots, 172
 - understanding schedules for volume sets, 166
 - using, 161
 - versus backups, 161
- SNMP
 - access control, 93
 - clients, adding, 93
 - configuring access control for SNMP clients, 93
 - disabling agents, 94
 - disabling traps, 95
 - editing trap recipients, 94
 - enabling agents, 92
 - enabling agents for event notification, 92
 - enabling traps, 94
 - overview, 92
 - removing trap recipient, 94
 - setting up for event notification, 92
 - using MIB, 95
 - using traps, 94
- software
 - upgrading storage systems, 25
- space allocation, 140
- space requirements
 - planning for SmartClone volumes, 179
- speed/duplex
 - configuring, 50
 - VSA, 48
- split network configuration, 48
- spoofing, 230
- starting
 - management group, 117
 - managers on storage systems, 114
 - virtual manager to recover quorum, 130
- static IP addresses and DNS, 71
- statistics
 - adding, 216
 - exporting performance to a CSV file, 220
 - in the Performance Monitor defined, 213
 - removing, 218
 - viewing details of, 217
- statistics sample data
 - clearing, 218
- status
 - dedicated boot devices, 29
 - NIC bond, 67
 - RAID, 36
 - safe to remove disk, 43
 - storage system, 137
 - storage system inoperable, 137
 - storage system overloaded, 137
- stopping
 - managers, 115
 - managers, implications of, 115
 - virtual manager, 131
- storage
 - adding to a cluster, 135
 - configuration on storage systems, 30
 - configuring, 30
 - overview, 30
 - provisioning, 140
 - upgrading in a cluster, 135
- storage pool, 103
- storage space
 - raw space, 151
- storage system inoperable, 137
- storage system overloaded, 137
- storage system status and VSA, 137
- storage systems
 - adding first one, 16, 104
 - adding to existing cluster, 135
 - adding to management group, 114
 - and raw space, 150
 - and space provisioned in cluster, 150
 - configuration categories, 21
 - configuration overview, 21
 - configuring, 17
 - details tab, 25
 - exchanging in a cluster, 136
 - finding on network, 13, 17, 19
 - ghost storage system, 138
 - locating in a rack, 23
 - locating in rack, 23
 - logging in to an additional, 22
 - not found, 20
 - powering off, 24
 - prerequisites for removing from management group, 118
 - rebooting, 24
 - registering, 28
 - removing, 119
 - removing from cluster, 136
 - removing from management group, 118
 - reorder position in a cluster, 136
 - repacing disks, 42
 - repairing in clusters, 138
 - saving statistics to a file, 98
 - statistics, 97
 - status of, 137
 - storage configuration of, 30
 - tasks, 21
 - upgrading software, 25, 26
- Subscriber's Choice, HP, 239
- support
 - setting up CMC for remote support, 20
- support bundles
 - exporting, 102
- support logs
 - exporting, 102
- support software, remote, 239

- swap
 - storage systems, 135
- synchronizing time on storage systems, 165
- system requirements
 - for Failover Manager on ESX Server, 122
- systems
 - finding on network, 13, 19

T

- TCP
 - frame size, 51
 - speed and duplex, 50
 - status, 50
 - status tab, 50
- TCP/IP tab, 53
- technical support
 - HP, 239
 - saving log files for, 100
 - service locator website, 240
- temporary space
 - deleting, 172
 - for read/write snapshots, 172
 - making application-managed snapshot available after converting, 171
- thresholds
 - changing for a snapshot, 165
 - requirements for changing in snapshots, 165
- time
 - editing NTP server, 77
 - NTP servers, preferred, not preferred, 77
 - selecting NTP, 76
 - setting with NTP, 76
 - setting without NTP, 78
 - synchronizing on storage systems, 165
 - zone, setting on storage system, 76, 78
- time remaining on evaluation period, 222
- time zone
 - changing for the Performance Monitor, 216
 - setting, 78
- toolbar
 - map view, 14
 - Performance Monitor window, 211
- total storage space in cluster, 148
- trap recipient
 - removing, 94
- traps
 - disabling SNMP, 95
 - editing SNMP recipient, 94
 - enabling SNMP, 94
 - sending test, 95
 - SNMP, 94
- troubleshooting clusters
 - repair storage system, 136
 - slow I/O, 136
- troubleshooting storage systems, finding, 19

U

- updating
 - manager IP addresses, 75

- upgrades
 - checking for available, 25
 - downloading, 25
 - setting preferences for, 25
- upgrading
 - available storage systems software, 26
 - SAN/iQ, 25, 26
 - storage system software, 25, 26
 - storage systems in clusters, 135
 - the CMC, 25, 26

- Use Summary window, 148

- user
 - adding a group to a user, 80
 - administrative, 79
 - administrative default user, 79
 - changing a user name, 79
 - deleting administrative, 80
 - editing, 80
 - password, 79

V

- verifying NIC bond, 66
- VI Client
 - recreating disk for VSA, 40
- viewing
 - clone points, volumes, and snapshots, 191
 - disk report, 38
 - disk setup report/disk, 37
 - RAID setup report, 32
 - SmartClone volumes, 189
 - statistics details, 217
- virtual IP address, 229
 - and iSCSI, 229
 - changing, iSCSI, 134
 - configuring for iSCSI for, 132
 - gateway session when using load balancing, 230
 - host storage system, 229
 - removing, iSCSI volume, 134
- virtual machine, 107
- virtual manager
 - adding, 129
 - configurations for using, 125
 - configuring, 129
 - function of, 125
 - overview, 125
 - removing, 131
 - starting to recover quorum, 130
 - stopping, 131
- virtual RAID
 - data safety and availability, 35
- virtual storage system
 - data protection, 33
 - RAID device, 33
- VMware ESX Server, 33
- VMware Server, 107
- volume access
 - Best Practice Summary, 113
- volume availability, 137
- volume sets

- creating application-managed snapshots for, [164](#)
- creating schedules to snapshots for, [166](#)
- deleting application-managed snapshots for, [159](#), [176](#)
- volume size
 - best practice for setting, [141](#)
- volume use window, [149](#)
- volume-level data protection
 - Best Practice Summary, [113](#)
- volumes
 - adding, [156](#)
 - application server requirements for rolling back, [173](#)
 - Assign Volume and Snapshot wizard, [18](#)
 - assigning to servers, [202](#), [203](#)
 - changing clusters, [159](#)
 - changing data protection levels, [159](#)
 - changing descriptions, [158](#)
 - changing size, [159](#)
 - comparing the load of two, [209](#)
 - consumed space, [149](#)
 - controlling server access to, [196](#)
 - creating SmartClone, [188](#)
 - creating using the Wizard, [17](#)
 - data protection level, [142](#)
 - deleting, [159](#)
 - editing, [157](#)
 - editing server assignments, [203](#), [204](#)
 - formatting for use, [204](#)
 - iSCSI, [230](#)
 - iSCSI and CHAP, [230](#)
 - logging on to, [204](#)
 - map view, [157](#)
 - mounting file systems on, [154](#)
 - overview, [154](#)
 - planning, [140](#), [154](#)
 - planning size, [140](#)
 - planning type, [154](#)
 - prerequisites for adding, [154](#)
 - prerequisites for deleting, [159](#), [160](#), [173](#), [176](#)
 - primary type, [154](#)
 - reclaimable space, [149](#)
 - remote type, [154](#)
 - requirements for adding, [154](#)
 - requirements for changing, [158](#)
 - restrictions on deleting, [159](#)
 - restrictions on rolling back, [173](#)
 - restriping, [46](#)
 - rolling back, [172](#)
 - saved space, [149](#)
 - setting as persistent targets, [204](#)
- volumes and snapshots
 - availability, [28](#)
- VSA
 - cloning, [103](#)
 - disk status, [40](#)
 - frame size, [48](#)
 - network interface, [48](#)
 - NIC bonding, [48](#)
 - NIC flow control, [48](#)
 - RAID rebuild rate, [35](#)

- reconfiguring RAID, [36](#)
- recreate disk, [40](#)
- speed/duplex, [48](#)
- storage system overloaded, [137](#)
- virtual RAID and data safety and availability, [35](#)

W

- warning events
 - defined, [85](#)
- warnings
 - all storage systems in a cluster operate at a capacity equal to that of the smallest capacity, [132](#)
 - changing RAID erases data, [36](#)
 - check Safe to Remove status, [43](#)
 - cloning VSA, [103](#)
 - deleting management group causes data loss, [119](#)
 - DHCP static IP addresses, [55](#)
 - DHCP unicast communication, [55](#)
 - return repaired system to same place, [139](#)
 - stopping manager can cause data loss, [115](#)
- websites
 - customer self repair, [240](#)
 - HP, [240](#)
 - HP Subscriber's Choice for Business, [239](#)
 - product manuals, [240](#)
- windows
 - alarms window, [13](#)
 - navigation window, [13](#)
 - tabs window, [13](#)
- write failure warnings, [141](#)