

HP LeftHand Storage Solutions

user guide



Legal and notice information

© Copyright 2009 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.

WARRANTY STATEMENT: To obtain a copy of the warranty for this product, see the warranty information website:

<http://www.hp.com/go/storagewarranty>

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

Contents

About this guide	27
Related documentation	27
HP technical support	27
Subscription service	27
HP websites	27
Documentation feedback	28
1 Getting started	29
Using the CMC	29
Auto discover	29
The CMC	29
Performing tasks in the CMC using the menu bar	30
Using the navigation window	31
Logging in	31
Traversing the navigation window	31
Single-clicking	31
Double-clicking	32
Right-clicking	32
Available nodes	32
CMC storage hierarchy	32
Icons	33
Using the tab window	33
Tab window conventions	33
Using the alert window	34
Setting naming conventions	34
Changing naming conventions	34
Creating storage by using the Getting Started Launch Pad	37
Prerequisites	37
Finding the storage nodes	37
Configuring storage nodes	37
Configure storage node categories	38
Creating a volume using the wizard	38
Enabling server access to volumes	39
Continuing with the SAN/iQ software	39
Finding storage nodes on an ongoing basis	39
Turn off Auto Discover for storage nodes	39
Troubleshooting—Storage nodes not found	39
Changing which storage nodes appear in the navigation window	40
Configuring multiple storage nodes	40
2 Working with storage nodes	43
Storage node configuration categories	43
Storage node configuration category definitions	43
Storage Node Tasks	43

Working with the storage node	44
Logging in to and out of storage nodes	44
Automatic log in	44
Logging out of a storage node	44
Changing the storage node hostname	44
Locating the storage node in a rack [NSM 260, DL 320s (NSM 2120), DL 380, and HP LeftHand P4500]	45
Backing up and restoring the storage node configuration	45
Backing up storage node does not save all settings	45
Backing up the storage node configuration file	46
Restoring the storage node configuration from a file	46
Powering off or rebooting the storage node	47
Powering on or off, or rebooting [NSM 4150]	47
Rebooting the storage node	48
Powering off the storage node	48
Upgrading the SAN/iQ software on the storage node	49
Prerequisites	49
Copying the upgrade files from web site	49
Upgrading the storage node	50
Registering advanced features for a storage node	50
Determining volume and snapshot availability	51
Checking status of dedicated boot devices	51
Checking boot device status	52
Getting there	52
Starting or stopping a dedicated boot device [NSM 160, NSM 260]	52
Powering on or rebooting storage nodes with two dedicated boot devices [NSM 160, NSM 260]	53
Replacing a dedicated boot device	53
NSM 160	53
NSM 4150	53
Replacing and activating a new boot flash card [NSM 160, NSM 260]	54

3 Storage 55

Configuring RAID and managing disks	55
RAID as a storage requirement	55
Getting there	55
Status indicators	56
Configuring and managing RAID	56
Benefits of RAID	57
RAID configurations defined	57
RAID0	57
RAID1/10	57
Storage Capacity in RAID1/10	57
RAID5, RAID5 + spare, or RAID50	58
Parity and storage capacity in RAID5 or 5 + spare	58
RAID5 and hot spare disks	58
RAID50 on the NSM 4150	59
RAID6	59
Parity and storage capacity in RAID6	59
Drive failure and hot swapping in RAID6	59
Explaining RAID devices in the RAID setup report	59
RAID devices by RAID type	60
Virtual RAID devices	60
Devices configured in RAID0	60

Devices configured in RAID10	61
Devices configured in RAID5	63
RAID5 in the NSM 260	65
RAID6 in the DL320s (NSM 2120), HP LeftHand P4500	66
RAID6 in the HP LeftHand P4300	66
Planning the RAID configuration	67
Data replication	67
Using RAID for data redundancy	67
Using volume replication in a cluster	68
Using RAID with replication in a cluster	68
Mixing RAID configurations	69
Setting RAID rebuild rate	69
Set RAID rebuild rate	70
Reconfiguring RAID	70
Requirements for reconfiguring RAID	70
Changing preconfigured RAID on a new storage node	70
Changing RAID on storage nodes in management groups	70
To reconfigure RAID	71
Monitoring RAID status	71
Data transfer and RAID status	71
Data redundancy and RAID status	71
Managing disks	72
Getting there	73
Reading the disk report on the Disk Setup tab	73
Verifying disk status	74
Viewing disk status for the NSM 160	74
Viewing disk status for the NSM 260	75
Viewing disk status for the DL380	75
Viewing disk status for the DL320s (NSM 2120)	76
Viewing disk status for the IBM x3650	77
Viewing disk status for the VSA	77
Viewing disk status for the Dell 2950 and NSM 2060	77
Viewing disk status for the NSM 4150	78
Viewing disk status for the HP LeftHand P4500	79
Viewing disk status for the HP LeftHand P4300	80
Replacing a disk	80
In the VSA	80
Overview of replacing a disk	81
Replacing disks in non-hot-swap platforms (IBM x3650)	81
Replacing disks in RAID0 configurations	82
Preparing for a disk replacement	82
To prepare for disk replacement	82
Identify physical location of storage node and disk	82
Best practice checklist for single disk replacement in RAID0	82
Best practice checklist for single disk replacement in RAID1/10, RAID5, RAID50, and RAID6	83
Replacing a disk in RAID0	83
Manually power off the disk in the CMC for RAID0	83
Physically replace the disk drive in the storage node	83
Manually powering on the disk in the CMC	83
Volume restriping	84
Replacing a disk in a non-hot-swap platform (IBM x3650)	84
Manually power off the disk in the CMC for RAID1/10 and RAID5	84
Physically replace the disk drive in the storage node	85
Manually powering on the disk in the CMC	85

RAID rebuilding	86
Replacing a disk in a hot-swap platform (NSM 160, NSM 260, DL380, DL320s [NSM 2120], Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300)	86
Replace the disk	86
Physically replace the disk drive in the storage node	86
RAID rebuilding	86
4 Managing the network	89
Network best practices	89
Changing network configurations	89
Best practices when changing network characteristics	89
Getting there	90
The TCP/IP tab	90
Identifying the network interfaces	90
Pinging an IP address	91
To ping an IP address	91
Configuring the IP address manually	91
Using DHCP	91
To set IP address using DHCP	92
Configuring network interface bonds	92
Best practices	93
NIC bonding and speed, duplex, frame size and flow control settings	93
How active-passive works	94
Physical and logical interfaces	94
Requirements for active-passive	95
Which physical interface is preferred	95
Summary of NIC status during failover	95
Example network configurations with active-passive	96
How link aggregation dynamic mode works	97
Requirements for link aggregation dynamic mode	97
Which physical interface is preferred	97
Which physical interface is active	97
Summary of NIC states during failover	98
Example network configurations with link aggregation dynamic mode	98
How adaptive load balancing works	98
Requirements for adaptive load balancing	98
Which physical interface is preferred	99
Which physical interface is active	99
Summary of NIC states during failover	99
Example network configurations with adaptive load balancing	100
Creating a NIC bond	100
Prerequisites	100
Bond guidelines	100
Creating the bond	101
Verify communication setting for new bond	102
Viewing the Status of a NIC Bond	103
Deleting a NIC bond	104
Verify communication setting after deleting a bond	105
Disabling a network interface	106
Disabling a network interface	106
If the storage node is in a management group	107
Configuring a disabled interface	107
TCP status	107
The TCP status tab	107

Managing settings on network interfaces	108
Requirements	108
Changing speed and duplex settings	108
Requirements	109
Best practice	109
To change the speed and duplex	109
Changing NIC frame size	109
Requirements	109
Best practices	110
Jumbo frames	110
Editing the NIC frame size	110
Changing NIC flow control	111
Requirements	111
Enabling NIC flow control	111
Using a DNS server	112
DNS and DHCP	112
DNS and static IP addresses	112
Getting there	112
Adding the DNS domain name	112
Adding the DNS server	112
Adding domain names to the DNS suffixes	113
Editing a DNS server	113
Editing a domain name in the DNS suffixes list	113
Removing a DNS server	114
Removing a domain suffix from the DNS suffixes list	114
Setting up routing	114
Adding routing information	114
Editing routing information	115
Deleting routing information	115
Configuring storage node communication	115
Selecting the interface used by the SAN/iQ software	116
Updating the list of manager IP addresses	117
Requirements	117

5 Setting the date and time 119

Management group time	119
Getting there	119
Refreshing the management group time	119
Using NTP	120
Editing NTP servers	120
Deleting an NTP server	121
Delete an NTP server	121
Changing the order of NTP servers	121
Editing the date and time	121
Editing the time zone only	122

6 Administrative users and groups 123

Getting there	123
Managing administrative users	123
Default administrative user	123
Adding a new administrative user	123
Editing administrative users	123
Changing a user's description	124
Changing a user's password	124

Adding group membership to a user	124
Removing group membership from a user	124
Deleting an administrative user	124
Managing administrative groups	125
Default administrative groups	125
Adding administrative groups	125
Editing administrative groups	126
Change the description of a group	126
Changing administrative group permissions	126
Adding users to an existing group	127
Removing users from a group	127
Deleting administrative groups	127

7 Using SNMP 129

Using SNMP	129
SNMP on the DL 380 and DL 320s (NSM 2120)	129
Getting there	129
Enabling SNMP agents	129
Community strings for DL 380 and DL 320s (NSM 2120)	130
Enabling an SNMP agent	130
Editing access control entries	131
Deleting access control entries	131
Using the SNMP MIB	132
Installing the LeftHand networks MIB	132
Disabling the SNMP agent	132
Disabling SNMP	132
Adding SNMP traps	133
Prerequisite	133
Enable SNMP traps	133
Removing trap recipients	133
Disabling SNMP traps	134

8 Reporting 135

Active monitoring overview	135
Using alerts for active monitoring	135
Getting there	136
Selecting alerts to monitor	136
Adding variables to monitor	136
Editing a monitored variable	137
Removing a variable from active monitoring	138
List of monitored variables	138
Setting CMC notification of alerts	141
Setting SNMP notification of alerts	141
Setting email notifications of alerts	141
Setting SMTP settings	142
Applying SMTP settings to a management group	142
Setting notification for one variable	142
Setting notification for several variables	143
Viewing and saving alerts	143
Saving the alert log of all variables	143
Saving the alert history of a specific variable	144
Using hardware information reports	144
Running diagnostic reports	144
Getting there	144

Viewing the diagnostic report	145
List of diagnostic tests	145
Using the hardware information report	150
Generating a hardware information report	150
Saving a hardware information report	151
Hardware information report details	152
Using hardware information log files	167
Saving log files	167
Using remote log files	168
Adding a remote log	168
Configuring the remote log target computer	169
Editing remote log targets	169
Deleting remote logs	169

9 Working with management groups 171

Functions of management groups	171
Requirements for creating a management group	171
Managers overview	171
Functions of managers	172
Managers and quorum	172
Regular managers and specialized managers	172
Failover managers	173
Virtual Managers	173
Creating a management group and default managers	173
Configuration summary overview	174
Summary roll-up	174
Configuration guidance	174
Best practices	175
Reading the configuration summary	175
Creating a management group	177
Guide to creating management groups	178
Getting there	178
Creating a new management group	178
Create management group and add storage nodes	179
Add administrative user	179
Set management group time	179
Create cluster and assign a VIP	179
Create a volume and finish creating management group	179
Adding a storage node to an existing management group	180
Logging in to a management group	180
Choosing which storage node to log in to	180
Logging out of a management group	181
Management group Maintenance tasks	181
Starting and stopping managers	181
Starting additional managers	181
Stopping managers	181
Editing a management group	182
Setting or changing the local bandwidth priority	183
Set or change local bandwidth priority	183
Backing up a management group configuration	183
Backing up a management group with remote copy relationships	184
Backup a management group configuration	184
Restoring a management group	184
Safely shutting down a management group	184

Prerequisites	185
Shut down the management group	185
If volumes are still connected to servers or hosts	185
Start the management group back up	185
Restarted management group in maintenance mode	186
Manually change management group to normal mode	186
Removing a storage node from a management group	187
Prerequisites	187
Remove the storage node	187
Deleting a management group	187
Prerequisites	187
Delete the management group	187
Setting the management group version	188

10 Using specialized managers 189

Definitions	189
Failover manager overview	189
Failover Manager requirements	189
Minimum system requirements for using with VMware server or player	189
Minimum system requirements for using with VMware ESX Server	190
Planning the virtual network configuration	190
Upgrading the 7.0 Failover Manager	190
Using Failover Manager with VMware Server or VMware Player	190
Installing and configuring the Failover Manager	190
Failover Manager configuration	190
To install the Failover Manager	191
Uninstalling the Failover Manager for VMware Server or Player	194
Troubleshooting the Failover Manager on VMware Server or Player	194
Fix startup/shutdown options	194
Fix network settings to find Failover Manager	195
Using the Failover Manager on VMware ESX Server	195
Installing the Failover Manager on VMware ESX Server	195
Using the HP LeftHand Management DVD	195
Using the HP LeftHand Networks web site download	196
For ESX 3.5+ or ESXi	196
For ESX Server 3.0 to 3.0.2	196
Configuring the Failover Manager using the VI Client	196
Add Failover Manager to inventory	196
Select a network connection	197
Power on the Failover Manager and configure the IP address and host name	197
Finishing up with VI Client	198
Uninstalling the Failover Manager from VMware ESX Server	199
Troubleshooting the Failover Manager on ESX Server	199
Virtual manager overview	200
When to use a virtual manager	200
Disaster recovery using a virtual manager	200
Management group across two sites with shared data	200
Management group in a single location with two storage nodes	200
Storage node maintenance using a virtual manager	201
Benefits of a virtual manager	201
Requirements for using a virtual manager	201
Configuring a cluster for disaster recovery	202
Best practice	202
Configuration steps	202

Adding a virtual manager	204
Starting a virtual manager to regain quorum	205
Starting a virtual manager	206
Verifying virtual manager status	206
Stopping a virtual manager	207
Removing a virtual manager	207
11 Working with clusters	209
Clusters and storage node capacity	209
Prerequisites	209
Creating additional clusters	209
Prerequisites	209
Number of storage nodes in clusters	209
To create additional clusters	209
Configure virtual IP and iSNS for iSCSI	210
New for release 8.0	210
Adding an iSNS server	210
Tracking cluster usage	210
Editing a cluster	211
Prerequisite	211
Getting there	211
Adding a new storage node to an existing cluster	211
Prerequisite	211
Adding storage to a cluster	212
Removing a storage node from a cluster	212
Changing or removing the virtual IP	212
Preparing servers	212
Changing the virtual IP address	213
Removing the virtual IP address	213
Finishing up	213
Changing or removing an iSNS server	213
Preparing clients	213
Changing an iSNS server	213
Deleting an iSNS server	213
Finishing up	214
Troubleshooting a cluster	214
Auto Performance Protection	214
Auto Performance Protection and the VSA	214
Auto Performance Protection and other clusters	215
Repairing a storage node	216
Prerequisites	216
Deleting a cluster	219
12 Provisioning storage	221
Understanding how the capacity of the SAN is used	221
Provisioning storage	221
Best practices	222
Provisioning volumes	222
Full provisioning	222
Thin provisioning	222
Best practice for setting volume size	222
Planning data replication	223
Replication level	223
How replication levels work	223

Replication priority	225
Best practice for setting replication levels and redundancy modes	225
Provisioning snapshots	226
Snapshots versus backups	226
The effect of snapshots on cluster space	226
Managing capacity using volume size and snapshots	226
How snapshots are created	226
Ongoing capacity management	227
Number of volumes and snapshots	227
Reviewing SAN capacity and usage	227
Cluster use summary	228
Volume use summary	230
Node use summary	232
Measuring disk capacity and volume size	233
Block systems and file systems	233
Storing file system data on a block system	234
Changing the volume size on the server	234
Increasing the volume size in Microsoft Windows	234
Increasing the volume size in other environments	235
Changing configuration characteristics to manage space	235
Snapshot temporary space	235
Managing snapshot temporary space	236

13 Using volumes 237

Volumes and server access	237
Prerequisites	237
Planning volumes	237
Planning how many volumes	237
Planning volume types	238
Guide for volumes	238
Creating a volume	239
Creating a basic volume	239
Configuring advanced volume settings [optional]	240
Configuring advanced volume settings	240
Editing a volume	240
To edit a volume	241
Changing the volume description	241
Changing the cluster	242
Changing the replication level	242
Changing the replication priority	242
Changing the size	242
Making an unavailable redundancy volume available	242
Deleting a volume	243
Prerequisites	243
New in release 8.0	243
To delete the volume	243

14 Using snapshots 245

Using snapshots overview	245
Snapshots versus backups	245
Prerequisites	245
Using snapshots	245
Single snapshots versus scheduled snapshots	246
Guide for snapshots	246

Planning snapshots	246
Source volumes for tape backups	247
Best practice	247
Data preservation before upgrading software	247
Best practice	247
Automated backups	247
Best practice	247
Planning how many snapshots	247
Creating a snapshot	248
Requirements for application-managed snapshots	248
Creating snapshots for volume sets	249
Editing a snapshot	250
Mounting or accessing a snapshot	250
Mounting the snapshot on a host	250
Making an application-managed snapshot available	250
Making an application-managed snapshot available on a stand-alone server	251
Making an application-managed snapshot available on a server in a Microsoft cluster	252
Managing snapshot temporary space	253
Convert the temporary space	253
Delete the temporary space	253
Creating a schedule to snapshot a volume	254
Best practices for scheduling snapshots of volumes	254
Creating schedules to snapshot a volume	255
Editing scheduled snapshots	255
Pausing and resuming scheduled snapshots	256
Pause a schedule	256
Resume a schedule	256
Deleting schedules to snapshot a volume	256
Scripting snapshots	257
Rolling back a volume to a snapshot or clone point	257
New in release 8.0	257
Requirements for rolling back a volume	258
Prerequisite	258
Rolling back a volume from a snapshot or clone point	258
Choosing a roll back strategy	259
Continue with standard roll back	259
Create a new SmartClone volume from the snapshot	259
Cancel the roll back operation	260
Deleting a snapshot	261
Prerequisites	261
Delete the snapshot	261

15 SmartClone volumes 263

Overview of SmartClone volumes	263
What are SmartClone volumes?	263
Prerequisites	263
Glossary	264
Example scenarios for using SmartClone volumes	264
Deploy multiple virtual or boot-from-SAN servers	264
Safely use production data for test, development and data mining	265
Clone a volume	265
Planning SmartClone volumes	265
Space requirements	265
Naming convention for SmartClone volumes	266

Naming and multiple identical disks in a server	266
Server access	267
Defining SmartClone volume characteristics	267
Naming SmartClone volumes	267
Shared versus individual characteristics	269
Clone point	272
Shared snapshot	274
Creating SmartClone volumes	276
To create a SmartClone volume	276
Viewing SmartClone volumes	279
Map view	279
Using views	280
Manipulating the Map View	280
Viewing clone points, volumes and snapshots	282
Viewing utilization of clone points and SmartClone volumes	282
Editing SmartClone volumes	283
To edit the SmartClone volumes	284
Deleting SmartClone volumes	284
Deleting the clone point	285
Deleting multiple SmartClone volumes	285
16 Working with scripting	287
Scripting documentation	287
17 Controlling server access to volumes	289
Former terminology (release 7.0 and earlier)	289
Adding server connections to management groups	290
Prerequisites	290
Editing server connections	291
Deleting server connections	292
Assigning server connections access to volumes	292
Assigning server connections from a volume	293
Assigning volumes from a server connection	293
Editing server connection and volume assignments	294
Editing server connection assignments from a volume	294
Editing server assignments from a server connection	294
Completing the iSCSI Initiator and disk setup	294
Persistent targets or favorite targets	295
HP LeftHand DSM for MPIO settings	295
Disk management	295
18 Monitoring performance	297
Prerequisites	297
Introduction to using performance information	297
What can I learn about my SAN?	298
Current SAN activities example	298
Workload characterization example	298
Fault isolation example	299
What can I learn about my volumes?	299
Most active volumes examples	299
Activity generated by a specific server example	300
Planning for SAN improvements	300
Network utilization to determine if NIC bonding could improve performance example	301

Load comparison of two clusters example	301
Load comparison of two volumes example	302
Accessing and understanding the Performance Monitor window	303
Performance Monitor toolbar	304
Performance monitoring graph	305
Performance monitoring table	306
Understanding the performance statistics	306
Monitoring and comparing multiple clusters	308
Performance monitoring and analysis concepts	309
Workloads	309
Access type	309
Access size	309
Access pattern	309
Queue depth	309
Changing the sample interval and time zone	309
Adding statistics	310
Viewing statistic details	312
Removing and clearing statistics	312
Removing a statistic	312
Clearing the sample data	312
Clearing the display	313
Resetting defaults	313
Pausing and restarting monitoring	313
Changing the graph	313
Hiding and showing the graph	314
Displaying or hiding a line	314
Changing the color or style of a line	314
Highlighting a line	315
Changing the scaling factor	315
Exporting data	315
Exporting statistics to a CSV file	315
Saving the graph to an image file	316

19 Registering advanced features 317

Evaluating advanced features	317
30—Day evaluation period	317
Tracking the time remaining in the evaluation period	318
Viewing licensing icons	318
Starting the evaluation period	318
Backing out of Remote Copy evaluation	318
Scripting evaluation	319
Turn on scripting evaluation	319
Turn off scripting evaluation	319
Registering advanced features	320
Using license keys	320
Registering available storage nodes for license keys	320
Submitting storage node feature keys	321
Entering license keys to storage nodes	321
Registering storage nodes in a management group	321
Submitting storage node feature keys	322
Entering license keys	323
Saving license key information	324
Saving and editing your customer information	324
Editing your customer information file	324

Saving your customer information	324
20 SNMP MIB information	325
The supported MIBs	325
21 Replacing disks appendix	327
Replacing disks and rebuilding data	327
Before you begin	327
Prerequisites	327
Replacing disks	328
Verify storage node not running a manager	328
Stopping a manager	328
Repair the storage node	328
Prerequisite	328
Replace the disk	329
In the NSM 160 or the NSM 260	329
In the DL380 or the IBM x3650	330
In the DL320s (NSM 2120), Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4300	330
Rebuilding data	330
Re-create the RAID array	330
Checking progress for RAID array to rebuild	331
Return storage node to cluster	331
Restarting a manager	332
Add repaired node to cluster	332
Rebuild volume data	333
Controlling server access	333
Change local bandwidth priority	333
Remove ghost storage node	333
Finishing up	334
22 iSCSI and the HP LeftHand Storage Solution	335
Number of iSCSI sessions	335
Virtual IP addresses	335
Requirements for using a virtual IP address	335
iSNS server	336
iSCSI load balancing	336
Requirements	336
Compliant iSCSI initiators	336
Authentication (CHAP)	337
Requirements for configuring CHAP	338
iSCSI and CHAP terminology	338
Sample iSCSI configurations	339
Best practice	340
About HP LeftHand DSM for MPIO	340
23 Using the Configuration Interface	341
Connecting to the Configuration Interface	341
Establishing a terminal emulation session on a Windows system	341
Establishing a terminal emulation session on a Linux/Unix system	341
Opening the Configuration Interface from the terminal emulation session	342
Logging in to the Configuration Interface	342
Configuring administrative users	343

Configuring a network connection	343
Deleting a NIC bond	344
Setting the TCP speed, duplex, and frame size	344
Removing a storage node from a management group	345
Resetting the storage node to factory defaults	345
24 Glossary	347
Terms Used	347
Index	353

Figures

1 Viewing the three parts of the CMC	30
2 Viewing the menu bar in the navigation window	31
3 Default naming conventions for snapshots and SmartClone volumes	34
4 Using the default naming for all the elements	36
5 The SAN/iQ software storage hierarchy	38
6 Storage node configuration categories	43
7 Disk enclosure not found as shown in Details tab	48
8 Confirming storage node power off	49
9 Availability tab	51
10 Viewing the storage configuration category for a storage node	56
11 Example of the capacity of disk pairs in RAID10	57
12 Parity distributed across a RAID5 set using four disks	58
13 Parity distributed across RAID6	59
14 RAID10 in an Dell 2950	60
15 RAID0 on an NSM 160	60
16 RAID0 on an NSM 260	61
17 RAID0 on a DL380	61
18 RAID0 on an IBM x3650	61
19 RAID10 on an NSM 160 (mirroring done at hardware level)	61
20 RAID1 on an NSM 260	61
21 RAID10 on a DL380	62
22 RAID10 on the DL320s (NSM 2120) and the HP LeftHand P4500	62
23 RAID1+0 in the HP LeftHand P4300	62
24 Initial RAID10 setup of the Dell 2950 and NSM 2060	62
25 Initial RAID10 setup of the NSM 4150	63
26 RAID5 set in an NSM 160	64
27 RAID5 + spare in an NSM 160	64
28 RAID5 set in a DL380	64
29 RAID5 set in the DL320s (NSM 2120), and the HP LeftHand P4500	64
30 RAID5 set in the HP LeftHand P4300	64
31 RAID5 set in a IBM x3650	64
32 Initial RAID5 setup of the Dell 2950 and NSM 2060	65

33	NSM 260 RAID5 using six-disk sets	65
34	NSM 260 RAID5 using five disks plus a hot spare	65
35	Initial RAID50 setup of the NSM 4150	66
36	DL320s (NSM 2120) and HP LeftHand P4500 RAID6 using two six-disk sets	66
37	Raid5 in P4300	67
38	Monitoring RAID status on the main CMC window	72
39	Example of columns in the Disk Setup tab	73
40	Viewing the Disk Setup tab in an NSM 160	74
41	Diagram of the drive bays in the NSM 160	75
42	Viewing the Disk Setup tab in an NSM 260	75
43	Diagram of the drive bays in the NSM 260	75
44	Viewing the Disk Setup tab in a DL380	76
45	Arrangement of drives in the DL380	76
46	Viewing the Disk Setup tab in a DL320s (NSM 2120)	76
47	Diagram of the drive bays in a DL320s (NSM 2120)	76
48	Viewing the Disk Setup tab in the IBM x3650	77
49	Arrangement of drives in the IBM x3650	77
50	Viewing the disk status of a VSA	77
51	Viewing the Disk Setup tab in a Dell 2950 or NSM 2060	78
52	Drive bays, with bezel on, in a Dell 2950 or NSM 2060	78
53	Drive bays, with bezel off, in a Dell 2950 or NSM 2060	78
54	Viewing the Disk Setup tab in a NSM 4150	78
55	Drive bays, with bezel on, in an NSM 4150	79
56	Drive bays, with bezel off, in an NSM 4150	79
57	Viewing the Disk Setup tab in a HP LeftHand P4500	79
58	Diagram of the drive bays in a HP LeftHand P4500	79
59	Viewing the Disk Setup tab in a HP LeftHand P4300	80
60	Diagram of the drive bays in a HP LeftHand P4300	80
61	Viewing a power off or missing disk	84
62	Viewing a power off or missing disk	85
63	RAID rebuilding on the RAID Setup tab	86
64	Disk rebuilding on the Disk Setup tab	87
65	Active-passive in a two-switch topology with server failover	96
66	Active-passive failover in a four-switch topology	96
67	Link aggregation dynamic mode in a single-switch topology	98
68	Adaptive Load Balancing in a two-switch topology	100
69	Searching for the bonded storage node on the network	102

70	Viewing a new active-passive bond	102
71	Verifying interface used for SAN/iQ communication	103
72	Viewing the status of an active-passive bond	103
73	Viewing the status of a link aggregation dynamic mode bond	104
74	Searching for the unbonded storage node on the network	105
75	Verifying interface used for SAN/iQ communication	106
76	Selecting the SAN/iQ software network interface and updating the list of managers	116
77	Viewing the list of manager IP addresses	117
78	Opening the hardware information window	151
79	Viewing the hardware information for a storage node	151
80	Failover manager in the available nodes pool	173
81	Virtual manager added to a management group	173
82	Configuration summary is created when the first management group is configured	174
83	Understanding the summary graphs	176
84	Warning when items in the management group are reaching safe limits	177
85	Error when some item in the management group has reached its limit	177
86	Manager quorum at risk	182
87	Notification of taking volumes offline	185
88	Manually setting management group to normal mode	186
89	VMware Console opens with Failover Manager installed and registered	192
90	Failover Manager boots up	192
91	Setting the host name and IP address	193
92	Confirming the new IP address	193
93	Logging in to the SAN/iQ Configuration Interface	197
94	Two-site failure scenarios that are correctly using a virtual manager	202
95	Adding storage nodes to cluster in alternating site order	204
96	2-way replicated volume on 2-site cluster	204
97	Management group with virtual manager added	205
98	Starting a virtual manager when storage node running a manager becomes unavailable	206
99	Checking the storage node status on the Details tab	215
100	Exchanging ghost storage node	218
101	Replacing the repaired storage node	218
102	Repaired storage node returns to proper place in cluster	219
103	Write patterns in 2-Way replication	224
104	Write patterns in 3-Way replication	224
105	Write patterns in 4-Way replication	224

106 Cluster tab view	228
107 Reviewing the Use Summary tab	228
108 Viewing the space saved or reclaimable in the Volume Use tab	231
109 Provisioned space shows temp space used	232
110 Viewing the Node Use tab	232
111 Stranded storage in the cluster	233
112 Viewing multiple volumes and snapshots	244
113 Deleting multiple volumes in one operation	244
114 Delete multiple snapshots from the volumes and snapshots node	257
115 Rolling back a volume	259
116 New volume with shared clone point	260
117 How SmartClone volumes, clone points and shared snapshots appear in the CMC	264
118 Duplicate names on duplicate datastores in ESX Server	266
119 Example of using a base name with 10 SmartClone volumes	268
120 Rename SmartClone volume from base name	269
121 Programming cluster with 10 SmartClone volumes, 1 clone point, and the source volume	270
122 Changing one SmartClone volume changes all associated volumes and snapshots	271
123 SysAdm cluster now has the 10 SmartClone volumes, 1 clone point, and the source volume	271
124 Navigation window with clone point	273
125 Clone point appears under each SmartClone volume	274
126 Navigation window with shared snapshots	275
127 Setting characteristics for SmartClone volumes	276
128 Creating multiple SmartClone volumes	278
129 Creating multiple SmartClone volumes	278
130 New SmartClone volumes in Navigation window	279
131 Viewing SmartClone volumes and snapshots as a tree in the Map View	280
132 Viewing the organic layout of SmartClone volumes and associated snapshots in the Map View	280
133 Toolbar with display tools in the Map View window	281
134 Using the Magnify tool with Map View tree	281
135 Highlighting all related clone points in navigation window	282
136 Clone point Details tab showing utilization graph	283
137 SmartClone volume Details tab showing utilization graph	283
138 Viewing volumes that depend on a clone point	285
139 List of SmartClone volumes in cluster	285
140 Server assignments in the navigation window and the Volumes and Snapshots tab	290

141	Warning after changing load balancing check box	292
142	Example showing overview of cluster activity	298
143	Example showing volume's type of workload	298
144	Example showing fault isolation	299
145	Example showing IOPS of two volumes	299
146	Example showing throughput of two volumes	300
147	Example showing activity generated by a specific server	300
148	Example showing network utilization of three storage nodes	301
149	Example comparing two clusters	302
150	Example comparing two volumes	302
151	Performance Monitor window and its parts	303
152	Performance Monitor toolbar	304
153	Example of a warning message	305
154	Performance Monitor graph	305
155	Performance Monitor table	306
156	Performance statistics and where they are measured	307
157	Add Statistics window	311
158	Edit Line window	314
159	Verifying the start of the 30-day evaluation period	317
160	Icons indicating license status for advanced features	318
161	Storage node with a license key	321
162	Registering advanced features for a management group	322
163	Selecting the feature key	322
164	Entering a license key	323
165	Viewing license keys	323
166	Warning if volumes are not replicated	329
167	Checking RAID rebuild status	331
168	Finding compliant initiator information	336
169	Viewing compliant iSCSI initiators	337
170	Differentiating types of CHAP	337
171	Viewing the MS iSCSI initiator to copy the initiator node name	339
172	Configuring iSCSI (shown in the MS iSCSI initiator) for a single host with CHAP	339
173	Adding an initiator secret for 2-way CHAP (shown in the MS iSCSI initiator)	340

Tables

1	Default names provided	34
2	Example of how default names work	35
3	Numbering conventions with no defaults enabled	36
4	Dedicated boot devices by storage node	51
5	Boot device status	52
6	RAID levels and default configurations for storage nodes	55
7	Status and color definitions	56
8	Storage capacity of RAID5 sets in storage nodes	58
9	Information in the RAID setup report	60
10	Data availability and safety in RAID configurations	68
11	Disk management tasks for storage nodes	73
12	Description of items on the disk report	73
13	Identifying the network interfaces on the storage node	90
14	Comparison of active-passive, link aggregation dynamic mode and adaptive load balancing bonding	93
15	Bonded network interfaces	94
16	NIC status in active-passive configuration	94
17	Example active-passive failover scenario and corresponding NIC status	95
18	NIC status during failover with Active-Passive	95
19	Link aggregation dynamic mode failover scenario and corresponding NIC status	97
20	NIC status during failover with link aggregation dynamic mode	98
21	Example adaptive load balancing failover scenario and corresponding NIC status	99
22	NIC status during failover with Adaptive Load Balancing	99
23	Status of and information about network interfaces	107
24	Setting storage node speed and duplex settings	108
25	Using default administrative groups	125
26	Descriptions of group permissions	126
27	Types of alerts for active monitoring	137
28	List of monitored variables	138
29	List of hardware diagnostic tests and pass / fail criteria for NSM 160 and NSM 260	146
30	List of hardware diagnostic tests and pass/fail criteria for DL 380, DL 320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300	147

31	List of hardware diagnostic tests and pass/fail criteria for IBM x3650	148
32	List of hardware diagnostic tests and pass/fail criteria for VSA	149
33	List of hardware diagnostic tests and pass/fail criteria for Dell 2950, NSM 2060, and NSM 4150	149
34	Selected details of the Hardware report for the NSM 160, NSM 260, and VSA	152
35	Selected details of the Hardware Report for DL 380, DL 320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300	157
36	Selected details of the hardware report for IBM x3650	162
37	Selected details of the hardware Report for Dell 2950, NSM 2060, and NSM 4150	163
38	Managers and quorum	172
39	Default number of managers added when a management group is created	174
40	Management group requirements	178
41	Guide to local bandwidth priority settings	183
42	Troubleshooting for ESX Server installation	199
43	Requirements for using a virtual manager	201
44	Recommended SAN configurations for provisioning storage	222
45	Volume provisioning methods	222
46	Setting a replication level for a volume	223
47	Storage node availability and volume access by replication level and priority setting	225
48	Information on the Use Summary tab	228
49	Information on the Volume Use tab	230
50	Information on the Node Use tab	232
51	Common native file systems	233
52	Characteristics for new volumes	238
53	Requirements for changing volume characteristics	240
54	Snapshot characteristics	246
55	Common applications' daily change rates	246
56	Requirements for scheduling snapshots	254
57	Characteristics for creating a schedule to snapshot a volume	255
58	Terms used for SmartClone features	264
59	Characteristics for new SmartClone volumes	267
60	Characteristics of SmartClone volumes	272
61	How it works - clone point	273
62	How it works - shared snapshots	275
63	Map View display tools	281
64	Requirements for changing SmartClone volume characteristics	283
65	Overview of configuring server access to volumes	289

66	Entering CHAP information in a new server	291
67	Server connection permission levels	293
68	Performance Monitor table columns	306
69	Performance Monitor statistics	307
70	Descriptions of advanced features	318
71	Safely backing out of Remote Copy evaluation	319
72	Safely backing out of scripting evaluation	320
73	Replacing the ghost storage node with the repaired storage node	332
74	Configuring iSCSI CHAP	338
75	iSCSI terminology	338
76	Logging in depends on where the storage node is	342
77	Identifying ethernet interfaces on the storage node	343
78	Glossary	347

About this guide

This guide provides information about:

- Configuring, managing, and maintaining the HP LeftHand Storage Solution. This guide encompasses hardware reporting configuration, the volume and snapshot features, and guidance for maintaining the SAN.

Related documentation

You can more information about HP's LeftHand products from the Manuals page of the HP Business Support Center website:

<http://www.hp.com/support/manuals>

In the Storage section, click **Disk Storage Systems > Left hand SANs** and then select your product.

HP technical support

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 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>

Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to storedocsFeedback@hp.com. All submissions become the property of HP.

1 Getting started

Welcome to the SAN/iQ software and the CMC. Use the CMC to configure and manage the HP LeftHand Storage Solution.

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

Using the CMC

Use the CMC to:

- Configure and manage storage nodes
- Create and manage clusters and volumes

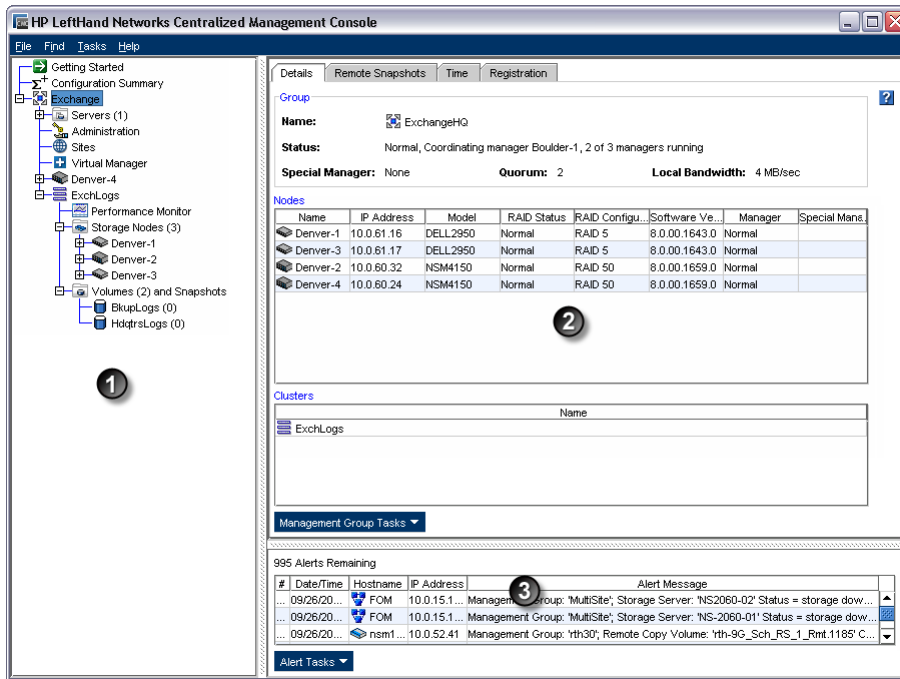
Auto discover

When you open the CMC the first time, it automatically searches the subnet for storage nodes. Any storage nodes it discovers appear in the navigation window on the left side of the CMC. If no storage nodes are found, the Find Nodes Wizard opens, and takes you through the steps to discover the storage nodes on your network.

Disable the Auto Discover feature by clearing the check box on the Find by Subnet and Mask window. For more information, see [“Finding the storage nodes”](#) on page 37.

The CMC

The CMC is divided into three sections.



1. Navigation window
2. Tab window
3. Alerts window

Figure 1 Viewing the three parts of the CMC

Navigation window—The left vertical pane displays the architecture of your network.

The physical and logical elements of your network include:

- Management groups
- Servers
- Administration
- Sites
- Failover Managers and Virtual Managers
- Clusters
- Storage Nodes and their configuration categories
- Volumes, including SmartClones
- Snapshots
- Remote Copies

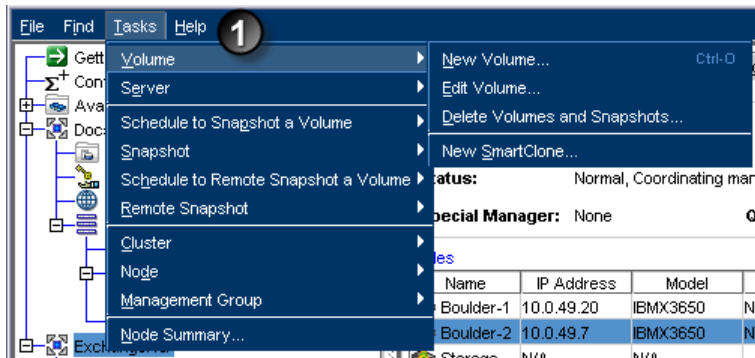
Tab window—For each element selected in the navigation window, the tab window on the right displays information about it. Commands related to the element are accessible from the Tasks menu on the bottom left of the tab window.

Alert window—View and delete alerts that display there.

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 nodes on the network 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 access online help and other information about the CMC and the SAN/iQ software.



1. Menu bar

Figure 2 Viewing the menu bar in the navigation window

Using the navigation window

The navigation window displays the components of your network architecture based on the criteria you set in the Find item in the menu bar, or by using the Find Storage Nodes wizard. Choose to display a small group of storage nodes, such as those in one management group, or display all storage nodes at one time.

Logging in

The CMC automatically logs in to storage nodes in the Available Nodes pool to access the configuration categories. After you have created management groups, and you then open the CMC, you must log in to access the management group. After you have logged in to one management group, the CMC attempts to log you in automatically to other management groups using the first login.

⚠ CAUTION:

Do not open the CMC on more than one machine at one time. Opening more than one session of the CMC on your network is not supported.

Traversing the navigation window

As you move through the items in the navigation window, the tab window changes to display the information and tabs about that item.

Single-clicking

Click an item once in the navigation window to select it.

Click the plus sign (+) once to open up a tree that displays more items.

Double-clicking

Double-click an item in the navigation window to open the hierarchy under that item. Double-click again to close it.

Right-clicking

Right-click an item in the navigation window to view a menu of commands for that item.

Getting started launch pad

The first item in the navigation window is always the Getting Started Launch Pad. Select the Launch Pad to access any of the three wizards to begin your work.

Available nodes

The second item in the navigation window is Available Nodes. Available Nodes includes the storage nodes and Failover Managers that are not in management groups. These storage nodes 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 [“Viewing the three parts of the CMC”](#) on page 30 .

CMC storage hierarchy

The items in the navigation window follow a specific hierarchy.

- **Management Groups**—Management groups are collections of storage nodes within which one or more storage nodes are designated as managers. Management groups are logical containers for the clustered storage nodes, volumes, and snapshots.
- **Servers**—Servers are application servers that you set up in a management group and assign to a volume to provide access to that volume.
- **Sites**—Sites are used to designate different geographical or logical sites in your environment. Sites are used with a Multi-Site SAN, and require a feature key. For more information about Multi-Site SANs, see the *HP LeftHand P4000 Multi-Site HA/DR Solution Pack User Manual* installed in the Documentation subdirectory with the CMC program files.
- **Clusters**—Clusters are groupings of storage nodes within a management group. Clusters contain the data volumes and snapshots.
- **Volumes**—Volumes store data and are presented to application servers as disks.
- **Snapshots**—Snapshots are copies of volumes. Snapshots can be created manually, as necessary, or scheduled to occur regularly. Snapshots of a volume can be stored on the volume itself, or on a different, remote, volume.
- **Remote Copies**—Remote copies are specialized snapshots that have been copied to a remote volume, usually at a different geographic location, using the SAN/iQ software feature, Remote Copy.

Icons

Each item in the navigation window has an icon depicting what type of item it is. A faded-looking icon indicates a remote item that is local or primary. A description is available of all the icons used in the CMC.

1. Click Help on the menu bar.
2. Select Graphical Legend from the menu.
3. View the Items tab and the Hardware tab.

The Items tab displays the icons that represent items, activities, and status in the navigation window.

The Hardware tab displays the icons that represent the different models of physical storage nodes that display in the navigation window.

Using the tab window

The tab window displays information about an item selected in the navigation window on the Details tab, as well as tabs for other functions related to that item. For example, “[Viewing the three parts of the CMC](#)” on page 30 shows the tabs that display when a management group is selected in the navigation window.

Tab window conventions

Tab windows have certain similarities:

- **Tabs**—Each tab in the tab window provides access to information and functions that relate to the element selected in the navigation window. For example, when a cluster is selected in the navigation window, the tabs contain information and functionality that relate specifically to clusters, such as usage information about volumes and storage nodes in that cluster and iSCSI sessions connected to the volumes.
- **Lists**—When presented with a list, such as a list of storage nodes as seen in a management group Details tab, you may select an item in the list to perform an action on.
- **Lists and right-click**—Right-click on an item in a list and a drop-down list of commands appropriate for that item appears.
- **Tasks buttons**—At the bottom of a tab window, the tasks button opens a menu of commands available for the element or function of that tab.

NOTE:

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

-
- **Sortable columns**—Click on a column head to sort the list in that column
 - **Sizable columns**—Drag a column boundary to the right or left to widen the column for better reading.

Using the alert window

Alert messages appear in the alert window as they are generated and are removed when the alert situation resolves in one of three ways:

- On its own
- When you remove them with the Alert Tasks commands
- When you close the CMC

To see old alerts, view those alerts from the storage node configuration categories, in the Alerts category.

- Select the Alert Log File tab and click the link to refresh the log file report. The log of alerts displays in the window.

Setting naming conventions

Use the Preferences window, opened from the Help menu, to set naming conventions for elements you create when building the HP LeftHand Storage Solution. Default values are provided, 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. Default names are disabled for management groups, clusters, and volumes.

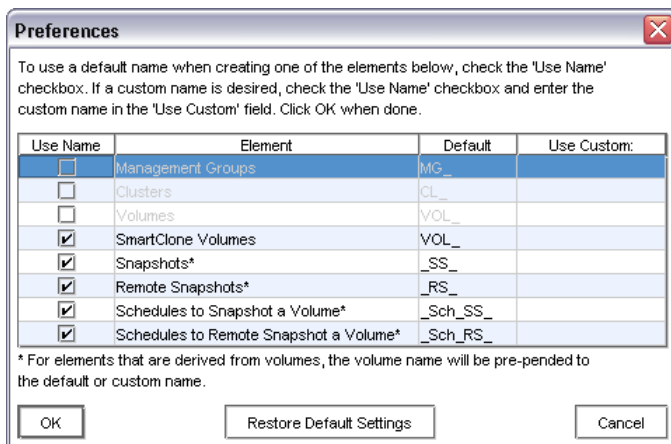


Figure 3 Default naming conventions for snapshots and SmartClone volumes

Changing naming conventions

Change the elements that use a default naming convention or change the naming convention itself.

Table 1 illustrates the default naming conventions built into the SAN/iQ software.

Table 1 Default names provided

Element	Default name
Disabled by default	
Management Groups	MG_

Element	Default name
Clusters	CL_
Volumes	VOL_
Enabled by default	
SmartClone Volumes	VOL_
Snapshots	_SS_
Remote Snapshots	_RS_
Schedules to Snapshot a Volume	_Sch_SS_
Schedules to Remote Snapshot a Volume	_Sch_RS_

If you were to use the given defaults for all the elements, the resulting names would look like those in the following example. 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
Disabled at installation		
Management Groups	MG_	MG_LogsBackup
Clusters	CL_	CL_OffSiteBkUp
Volumes	VOL_	VOL_DailyBkUp
Enabled at installation		
SmartClone Volumes	VOL_	VOL_VOL_DailyBkUp_SS_3_1
Snapshots	_SS_	VOL_DailyBkUp_SS_1
Remote Snapshots	_RS_	VOL_DailyBkUp_RS_1
Schedules to Snapshot a Volume	_Sch_SS_	VOL_DailyBkUp_Sch_SS_1.1
Schedules to Remote Snapshot a Volume	_Sch_RS_	VOL_DailyBkUp_Sch_RS_1_Pri.1VOL_DailyBkUp_Sch_1_RS_Rmt.1

This example is illustrated in [Figure 4](#).

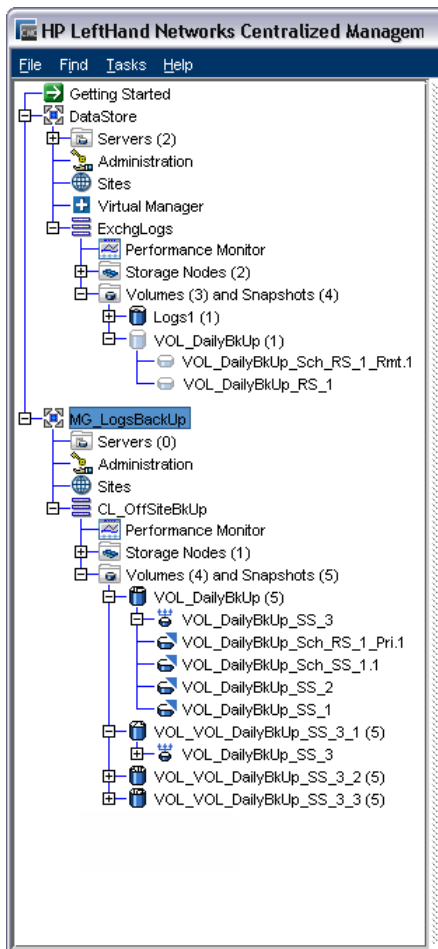


Figure 4 Using the default naming for all the elements

If you do not use any of the default names, then the only automatically generated naming elements are those that incrementally number a series of snapshots or SmartClone volumes. See [Table 3](#) .

Table 3 Numbering conventions with no defaults enabled

Element	Default name
Disabled at installation	
Management Groups	None
Clusters	None
Volumes	None
Enabled at installation	
SmartClone Volumes	<i>Name_#</i>
Snapshots	None
Remote Snapshots	None
Schedules to Snapshot a Volume	<i>Name.#</i>

Element	Default name
Schedules to Remote Snapshot a Volume	Name_Pri.#Name_Rmt.#

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, you will work through these steps with one storage node, 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 nodes on your network.
- Know the IP address or host name you configured with the KVM or serial Configuration Interface when you installed the storage node.
- Install the HP LeftHand Centralized Management Console software on a management workstation or server that can connect to the storage nodes on the network.
- Install an iSCSI initiator on the application server(s), such as the latest version of the Microsoft iSCSI initiator.



NOTE:

The HP LeftHand DSM for MPIO is the only supported multi-path solution for the HP LeftHand Storage Solution. Starting with SAN/iQ software release 7.0, you must install the Microsoft MPIO DSM if you want to use the HP LeftHand DSM for MPIO

Finding the storage nodes

Open the CMC, and using the Getting Started Launch Pad, start the Find Nodes Wizard.

To use the wizard, you need to know either the

- The subnet and mask of your storage network or
- The IP addresses or host names of the storage nodes

When you have found the storage nodes, they appear in the Available Nodes pool in the navigation window.

Configuring storage nodes

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

The most important categories to configure are:

- **RAID**—The storage node is shipped with RAID already configured and operational. Find instructions for ensuring that drives in the storage node are properly configured and operating in [Chapter 3](#) on page 55.
- **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 [Chapter 4](#) on page 89.

- Alerts - Use the email alert function or SNMP to ensure you have immediate access to up-to-date alert and reporting information. Detailed information for setting up SNMP and alerts can be found in [Chapter 7](#) on page 129 and “[Using alerts for active monitoring](#)” on page 135.

Configure storage node categories

1. From the navigation window, select a storage node in the Available Nodes pool.
2. Double-click to open the tree underneath the storage node.
The list of storage node configuration categories opens.
3. Select the Storage category.
The Storage tab window opens.
4. Select the RAID Setup tab and verify the RAID settings.
5. In the list of configuration categories, select the TCP/IP Network category and configure your network settings.
6. In the list of configuration categories, select the SNMP and/or Alerts categories to configure the monitoring for your IP SAN.

Creating a volume using the wizard

Next, you create the storage hierarchy 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 first task in the wizard is to assign one or more storage nodes to a management group. The second task is to cluster the storage nodes. The third task is to create a storage volume. This storage hierarchy is depicted in [Figure 5](#).

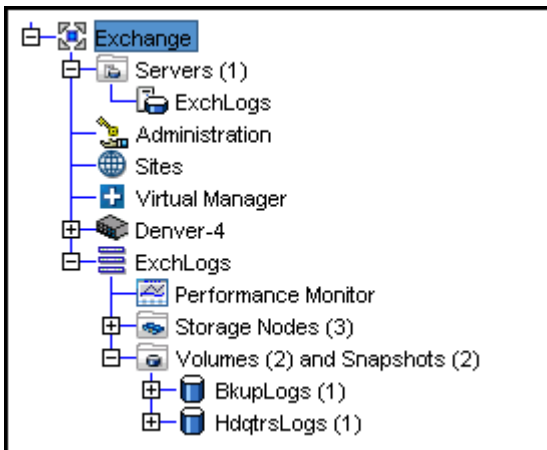


Figure 5 The SAN/iQ software storage hierarchy

While working through the wizard, you need to have ready the following information:

- A name for your management group (note: this cannot be changed in the future without destroying the management group)
- A storage node that you identified with the Find wizard and then configured
- A name for the cluster
- A name for the volume

- The size of the volume

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 [Chapter 17](#) on page 289 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 node name, and CHAP information if you plan to use CHAP.

Continuing with the SAN/iQ software

This section describes techniques for working with the CMC on an ongoing basis. It also describes how to copy the configuration of one storage node to other storage nodes.

Finding storage nodes on an ongoing basis

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 nodes that are found.

Turn off Auto Discover for storage nodes

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

1. From the menu bar, select Find > By Subnet and Mask.
2. Clear the Auto Discover check box.

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

However, if you have a subnet and mask listed, it will continue to search that subnet for storage nodes.

Troubleshooting—Storage nodes not found

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

1. If the storage node you are looking for does not appear in the navigation window, search again using the Find menu.
2. If you have searched by Subnet and Mask, try using the Find by IP or Host Name search or vice versa.
3. If searching again does not work, try the following:
 - Check the physical connection of the storage node.
 - Wait a few minutes and try the search again. If activity to the storage node was frequent, the storage node might not have responded to the search.

Possible reasons for not finding storage nodes

Other problems can occur that prevent CMC from finding a storage node:

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

Changing which storage nodes appear in the navigation window

1. Click the Find menu.
2. Select Clear All Found Items to remove all storage nodes from the navigation window.
3. Perform a Find using either method—By Subnet and Mask or By Node IP or Host Name—to find the desired set of storage nodes.

NOTE:

Control which storage nodes appear in the navigation window by entering only specific IPs or Host Names in the IP and Host Name List window. Then, when you open the CMC, only those IPs or Host Names will appear in the navigation window. Use this method to control which management groups appear.

Configuring multiple storage nodes

After you have configured one storage node with settings for alerts, SNMP monitoring, and remote log files, you can copy those settings between storage nodes.

For information about configuring these settings, see the following sections.

- [“Enabling SNMP agents”](#) on page 129
- [“Using alerts for active monitoring”](#) on page 135
- [“Using remote log files”](#) on page 168
- [“Setting email notifications of alerts”](#) on page 141

CAUTION:

Copying the configuration between different models may result in a monitored variable configuration with unsupported variables, incorrect thresholds, or removed variables. Be sure to verify that the configuration is correct on the storage nodes you copy to.

1. In the navigation window, select the storage node that has the configuration that you want to copy to other storage nodes.
2. Click Storage Node Tasks on the Details tab and select Copy Configuration.

1. In the Configuration Settings section, select which configurations you want to copy.
2. In the Copy Configurations to nodes section, select the storage nodes to which you want to copy the configurations.
3. Click Copy.
The configuration settings are copied to the selected storage nodes.
4. Click OK to confirm the operation and close the window.

2 Working with storage nodes

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

- Alerts
- Hardware
- SNMP
- Storage
- TCP/IP Network

Storage node configuration categories

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




Figure 6 Storage node configuration categories

Storage node configuration category definitions

The storage node configuration categories are described below.

- **Alerts**—Configure active monitoring settings of selected monitored variables and notification methods for receiving alerts.
- **Hardware**—Use the hardware category to run hardware diagnostic tests, to view current hardware status and configuration information, and to save log files.
- **SNMP**—Monitor the storage node using an SNMP management station. You can also enable SNMP traps.
- **Storage**—Manage RAID and the individual disks in the storage node.
- **TCP/IP Network**—For each storage node, configure and manage the network settings, including network interface cards (NICs), DNS servers, the routing table, and which interface carries SAN/iQ communication.

Storage Node Tasks

This section describes how to perform basic storage node tasks:

- [“Working with the storage node”](#) on page 44
- [“Logging out of a storage node”](#) on page 44
- [“Changing the storage node hostname”](#) on page 44
- [“Locating the storage node in a rack \[NSM 260, DL 320s \(NSM 2120\), DL 380, and HP LeftHand P4500\]”](#) on page 45
- [“Backing up and restoring the storage node configuration”](#) on page 45
- [“Rebooting the storage node”](#) on page 48
- [“Powering off the storage node”](#) on page 48

Working with the storage node

After finding all the storage nodes on the network, you configure each storage node individually.

1. Select the storage node in the navigation window.

Usually you will be logged in automatically. However, you will have to log in manually for any storage nodes running a software version earlier than release 7.0. If you do need to manually log in, the Log In window opens.

2. Type a user name and password.
3. Click Log In.

Logging in to and out of storage nodes

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 nodes in that group. You can log out of individual storage nodes in the management group, and log back in to them individually.

Automatic log in

Once you have logged into 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. Type the correct user name and password.
2. Click Log In.

Logging out of a storage node

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



NOTE:

If you are logged in to multiple storage nodes, you must log out of each storage node individually.

Changing the storage node hostname

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

1. In the navigation window, log in to the storage node.
2. On the Details tab, click Storage Node Tasks and select Edit Hostname.
3. Type 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 node in a rack [NSM 260, DL 320s (NSM 2120), DL 380, and HP LeftHand P4500]

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

1. Select a storage node in the navigation window and log in.
2. Click Storage Node Tasks on the Details tab and select Set ID LED On.

The ID LED on the front of the storage node illuminates a bright blue. Another ID LED is located on the back of the storage node.

When you click Set ID LED On, the status changes to On.

3. Select Set ID LED Off when you are finished.

The LED on the storage node turns off.

Backing up and restoring the storage node configuration

Back up and restore storage node configurations to save the storage node configuration to a file for use in case of a storage node failure. When you back up a storage node configuration, the configuration information about the storage node is stored in a file. If a storage node failure occurs, restore the backed up configuration to a replacement storage node. The replacement storage node will be configured identically to the original storage node at the time it was backed up.

 **NOTE:**

You must restore the configuration to the replacement storage node BEFORE you add it to the management group and cluster.

Backing up storage node does not save all settings

Backing up the configuration file for a storage node does not save data. Neither does it save information about the configuration of any management groups or clusters that the storage node belongs to. It also does not back up license key entries for registered features.

- To save the management group configuration, see “[Backing up a management group configuration](#)” on page 183.
- To preserve a record of the management group license keys, see “[Saving license key information](#)” on page 324.

 **NOTE:**

Back up the storage node configuration every time you change storage node settings. This ensures that you can restore a storage node to its most recent configuration.

Manual configuration steps following the restore

After you restore the storage node configuration from a file, up to three manual configuration steps are required:

- You must manually configure RAID on the storage node.
- You must manually add network routes after the restoration. Restoring a configuration file from one storage node to a second storage node does not restore network routes that were configured on the storage node.
- If you restore multiple storage nodes from one configuration file, you must manually change the IP address on the subsequent storage nodes. For example, if you back up the configuration of a storage node with a static IP address, and you then restore that configuration to a second storage node, the second storage node will have the same IP address.

Backing up the storage node configuration file

Use Back Up to save the storage node configuration file to a directory you choose.

1. In the navigation window, select a storage node.
2. Click Storage Node Tasks on the Details tab and select Back Up or Restore.
3. Click Back Up.
4. Navigate to a folder to contain the storage node configuration backup file.
5. Enter a meaningful name for the backup file or accept the default name (Storage Node_Configuration_Backup).

 **NOTE:**

The configuration files for all storage nodes that you back up are stored in the location you choose. If you back up multiple storage nodes to the same location, be sure to give each storage node configuration file a unique and descriptive name. This makes it easier to locate the correct configuration file if you need to restore the configuration of a specific storage node.

6. Click Save.

Restoring the storage node configuration from a file

Before you add the replacement storage node to the management group and cluster, use the configuration backup file to restore the configuration of the failed storage node to the replacement node. You may also need to manually configure RAID, network routes and, if you apply the same configuration backup file to more than one storage node, a unique IP address. You must also complete the manual configuration before adding the replacement storage node to the management group and cluster.

1. In the navigation window, select the storage node from the Available Nodes pool.

2. Click Storage Node Tasks on the Details tab and select Back Up or Restore.
3. Click Restore.
4. In the table, select the storage node you want to restore.
You may select multiple storage nodes to restore from the table.
5. Select the radio button: Install file on selected storage nodes one at a time (Recommended).
6. Click Browse to navigate to the folder where the configuration backup file is saved.
7. Select the file to restore and click Open Backup File.
8. Review the version and description to be sure that you are restoring the correct file.
9. Click Install.
10. When the restoration is complete, the Save to File and Close buttons on the Install Status window become enabled.
To save a log file of the restore operation before rebooting, click Save to File.
11. Click Close to finish restoring the configuration.
The storage node reboots and the configuration is restored to the identical configuration as that of the backup file.
12. Complete the configuration of the replacement storage node by reconfiguring the following characteristics as described in [“Manual configuration steps following the restore”](#) on page 46:
 - RAID
 - Network routes
 - IP address

Powering off or rebooting the storage node

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

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

Powering off an individual storage node is appropriate for servicing or moving that storage node. However, if you want to shut down more than one storage node in a management group, you should shut down the management group instead of individually powering off the storage nodes in that group. See [“Safely shutting down a management group”](#) on page 184.

Powering on or off, or rebooting [NSM 4150]

When powering on the NSM 4150, be sure to power on the two components in the following order:

1. Disk enclosure.
2. System controller.

Allow up to six minutes for the system controller to come up completely and be discovered by the CMC. If you cannot discover the NSM 4150 using the CMC after six minutes, contact Customer Support.

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




Figure 7 Disk enclosure not found as shown in Details tab

When powering off the NSM 4150, be sure to power off the two components in the following order:

1. Power off the system controller from the CMC as described in “Powering off the storage node” on page 48.
2. Manually power off the disk enclosure.

When you reboot the NSM 4150, use the CMC, as described in “Rebooting the storage node” on page 48. This process reboots only the system controller.

Rebooting the storage node

1. Select a storage node in the navigation window and log in.
2. Click Storage Node Tasks on the Details tab and select Power Off or Reboot.
3. In the minutes field, type 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 node reboots shortly after you complete [Step 5](#).



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 node starts the reboot in the specified number of minutes. The reboot takes several minutes.
6. Search for the storage node to reconnect the CMC to the storage node once it has finished rebooting.
See “[Finding the storage nodes](#)” on page 37.

Powering off the storage node

1. Log in to the storage node.
2. Select Storage Node 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, type 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 node powers off shortly after you complete [Step 5](#).

 **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.

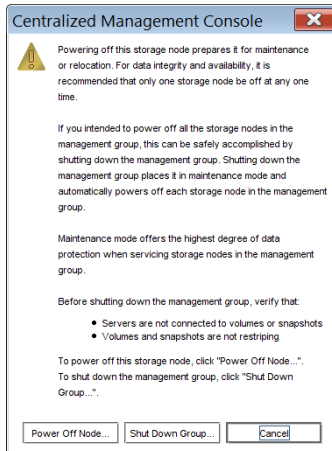


Figure 8 Confirming storage node power off

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 node

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

Prerequisites

Stop any applications that are accessing volumes that reside on the storage node you are upgrading and log off all related iSCSI sessions.

To view a list of available upgrades, select Check for Upgrades from the Help menu.

Copying the upgrade files from web site

Upgrade the SAN/iQ software on the storage node when an upgrade or a patch is released. The SAN/iQ software upgrade/installation takes about 10 to 15 minutes (may be longer on certain platforms), including the storage node reboot.

 **NOTE:**

For those models that contain 2 boot flash cards, both boot flash cards must be in place to upgrade the SAN/iQ software. See “[Checking status of dedicated boot devices](#)” on page 51.

Upgrading the storage node

Install upgrades on storage nodes individually, which is recommended. If you are upgrading multiple storage nodes that are not in a management group, you can upgrade them simultaneously.

 **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.

1. Log in to the first storage node you want to upgrade.
2. Click Storage Node Tasks on the Details tab and select Install Software.
3. From the list, select the storage node that you want to upgrade. Use the CTRL key to select multiple storage nodes to upgrade from the list.
4. Select this radio button: Install file on selected storage nodes one at a time (Recommended).
5. Click Browse to navigate to the folder where you copied the upgrade or patch file.
6. Select the file and click Open Install File.

Focus returns to the Install Software window. When the file name is present, the Install button becomes enabled.

7. Review the version and description to be sure that you are using the correct upgrade file.
8. Click Install.

Select the check box to have the install messages automatically scroll. These messages can be saved to a file.

(Optional) After the installation completes, click Save To File and choose a name and location for the file.

After the installation completes, the system reboots. After the system comes back online, it conducts a post-install qualification. After the system passes the post-install qualification, the upgrade process is complete.

9. Click Close when the installation is completed.

Registering advanced features for a storage node

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

For more information about registering advanced features, see [Chapter 19](#) on page 317.

Determining volume and snapshot availability

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

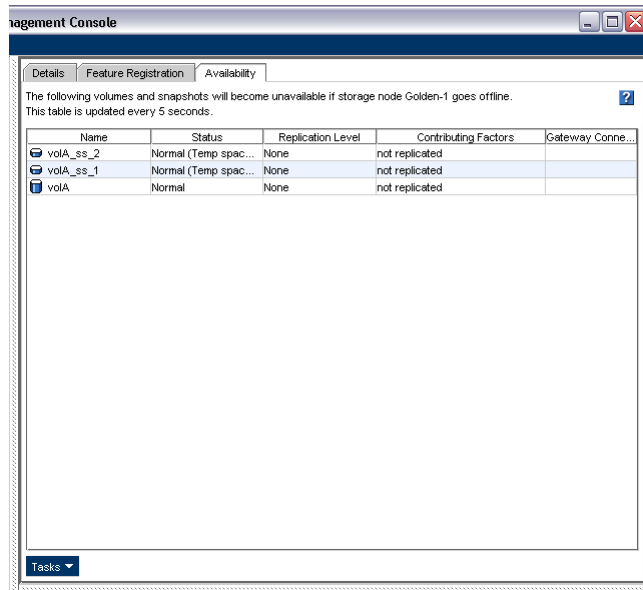


Figure 9 Availability tab

Checking status of dedicated boot devices

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

Table 4 Dedicated boot devices by storage node

Platform	Number and type of boot devices
NSM 160	2 compact flash cards
NSM 260	1 compact flash card
NSM 4150	2 hard drives

In storage nodes 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 LeftHand Networks Technical Support.

The following storage nodes do not have dedicated boot devices:

- DL 380
- DL 320s (NSM 2120)
- IBM x3650
- VSA

- Dell 2950
- NSM 2060
- HP LeftHand P4500

Checking boot device status

View dedicated boot device status in the Boot Devices tab window in the Storage category in the storage node tree. The platforms that have dedicated boot devices are listed in [Table 4](#) on page 51.

Getting there

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

The status of each dedicated boot device on the storage node is listed in the Status column. [Table 5](#) on page 52 describes the possible status for boot devices.



NOTE:

Some statuses only occur in a storage node with two boot devices.

Table 5 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 node. It will not be used to boot the storage node.
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 node. 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 alert is generated. See [“Using alerts for active monitoring”](#) on page 135.

Starting or stopping a dedicated boot device [NSM 160, NSM 260]

On storage nodes with dedicated boot devices, use this procedure to remove a boot device from service and later return it.

1. From the navigation window, select the storage node and log in if necessary.

2. Open the tree below the storage node and select Storage.
3. Select the Boot Devices tab.
4. In the tab window, select the boot device you want to start or stop.
5. Click Boot Devices Tasks and select either
 - Activate to make a boot device available in the event of a failure, or
 - Deactivate to remove a boot device from service.

Powering on or rebooting storage nodes with two dedicated boot devices [NSM 160, NSM 260]

When a storage node with two dedicated boot devices powers on or reboots, it references boot configuration information from one of two compact flash cards, located on the front of the storage node.

The storage node boot configuration information is mirrored between the two compact flash cards. If one card fails or is removed, the system can still boot. If you remove and replace one of the cards, you must activate the card to synchronize it with the other card.

NOTE:

There must always be at least one active flash card in the storage node. If you are upgrading the SAN/iQ software, a dual boot device storage node must contain both flash cards.

Replacing a dedicated boot device

NSM 160

If a compact flash card fails, first try to activate it on the Boot Devices window. If the card fails repeatedly, replace it with a new one.

You can also replace a boot flash card if you have removed the original card to store it as a backup in a remote location.

CAUTION:

A flash card from one storage node cannot be used in a different storage node. If a card fails, replace it with a new flash card.

NSM 4150

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

Removing a boot flash card [NSM 160, NSM 260]

Before you remove one of the boot flash cards from the storage node, deactivate the card in the CMC.

1. On the Boot Devices window, select the flash card that you want to remove.

2. Click Deactivate.

The flash card status changes to Inactive. It is now safe to remove the card from the storage node.

3. Power off the storage node.
4. Remove the flash card from the front of the storage node.

Replacing and activating a new boot flash card [NSM 160, NSM 260]

If you replace a boot flash card in the storage node, you must activate the card before it can be used. Activating the card erases any existing data on the card and then synchronizes it with the other card in the storage node.

1. Insert the new flash card in the front of the storage node.
2. Power on the storage node.
3. Log in to the storage node.
4. On the Boot Devices window, select the new flash card.
5. Click Activate.

The flash card begins synchronizing with the other card. When synchronization is complete, 'Active' displays in the Status column.

3 Storage

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

Configuring RAID and managing disks

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

RAID as a storage requirement

RAID must be configured for data storage. HP LeftHand Networks physical storage nodes come with RAID preconfigured. The VSA comes with RAID preconfigured, if you have first configured the data disk in the VI Client, as described in the *VSA Quick Start Guide*. The descriptions of RAID levels and configurations for the various storage nodes are listed in [Table 6](#).

Table 6 RAID levels and default configurations for storage nodes

Model	Preconfigured for	Available RAID levels
NSM 160	RAID5	0, 10, 5, 5 + spare
NSM 260	RAID5	0, 1, 5, 5 + spare
DL380	RAID5	0, 10, 5
DL320s (NSM 2120)	RAID5	10, 5, 6
IBM x3650	RAID5	0, 10, 5
Dell 2950	RAID10 or RAID5	10, 5 (cannot change RAID level)
NSM 2060	RAID10 or RAID5	10, 5 (cannot change RAID level)
HP LeftHand P4300	RAID5	5, 6, 10
NSM 4150	50	10, 50
VSA	Virtual RAID (if data disk is configured in the VI Client first)	RAID (virtual)
HP LeftHand P4500	RAID5	10, 5, 6

Getting there

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

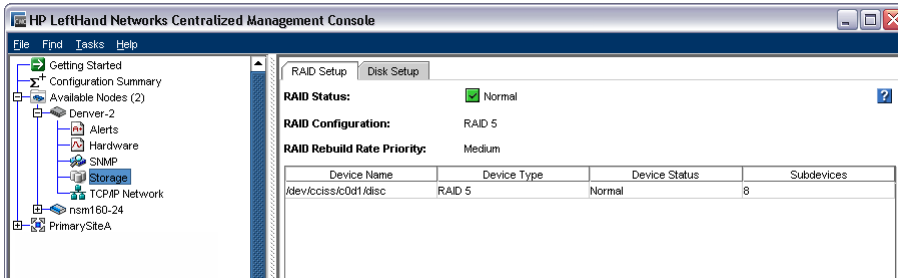


Figure 10 Viewing the storage configuration category for a storage node

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](#) 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
Hot Spare	Green
Hot Spare Down	Yellow

Configuring and managing RAID

Managing the RAID settings of a storage node 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 node
- Reconfiguring RAID when necessary

Benefits of RAID

RAID combines several physical disks into a larger logical disk. This larger logical disk can be configured to improve both read and write performance and data reliability for the storage node.

RAID configurations defined

The RAID configuration you choose depends upon how you plan to use the storage node. The storage node can be reconfigured with RAID0, RAID1 /10, RAID5, RAID5+hot spare, RAID50, or RAID6, depending on the model. See [Table 6](#) on page 55 for a list of RAID levels by model.

RAID0

RAID0 creates a striped disk set. Data is stored across all disks in the array, which increases performance. However, RAID0 does not provide fault tolerance. If one disk in the RAID set fails, all data on the set is lost.

Storage node capacity in RAID0 is equal to the sum total capacity of all disks in the storage node.

RAID1/10

RAID1 /10 combines mirroring data within pairs of disks and striping data across the disk pairs. RAID1 /10 combines data redundancy or disk mirroring (RAID1) with the performance boost of striping (RAID0).

Storage Capacity in RAID1/10

Storage capacity in RAID1/10 is half the total capacity of RAID0 in the storage node. The capacity of a single disk pair is equal to the capacity of one of the disks, thus yielding half the total capacity. Or, to put it another way,

RAID10 capacity = (single disk capacity x total # of disks) / 2




Figure 11 Example of the capacity of disk pairs in RAID10

RAID5, RAID5 + spare, or RAID50

RAID5 provides data redundancy by distributing data blocks across all disks in a RAID set. Redundant information is stored as parity distributed across the disks. The figure shows an example of the distribution of parity across four disks in a RAID5 set.




Figure 12 Parity distributed across a RAID5 set using four disks

Parity allows the storage node to yield more disk capacity for data storage than RAID10 allows.

Parity and storage capacity in RAID5 or 5 + spare

Parity in a RAID5 set equals the capacity of one disk in the set. Therefore, the capacity of any RAID5 set is $n - 1$, as illustrated in [Table 8](#).

Table 8 Storage capacity of RAID5 sets in storage nodes

Model	Number of disks in RAID5 set	Storage capacity of disks
NSM 160	4 disks 3 disks plus a spare	3 x single disk capacity 2 x single disk capacity
NSM 260	5 disks plus a spare (x 2 RAID sets) 6 disks (x 2 RAID sets)	8 x single disk capacity 10 x single disk capacity
DL380	6 disks	5 x single disk capacity
DL320s (NSM 2120)	6 disks (x 2 RAID sets)	10 x single disk capacity
IBM x3650	6 disks	5 x single disk capacity
Dell 2950	6 disks	5 x single disk capacity
NSM 2060	6 disks	5 x single disk capacity
HP LeftHand P4300	8 disks	7 x single disk capacity
NSM 4150	5 disks	4 x single disk capacity
HP LeftHand P4500	6 disks (x 2 RAID sets)	10 x single disk capacity

RAID5 and hot spare disks

RAID5 configurations that use a spare designate as a hot spare the remaining disk of the RAID set. With a hot spare disk, if any one of the disks in the RAID5 set fails, the hot spare disk is automatically added to the set and RAID starts rebuilding.

Table 8 on page 58 lists the RAID5 configurations by model, and indicates which configurations support a hot spare.

RAID50 on the NSM 4150

RAID50 combines distributing data blocks across disks in a RAID5 set and striping data across multiple RAID5 sets. RAID50 combines data redundancy (RAID5) with the performance boost of striping (RAID0).

The total capacity of the NSM 4150 in RAID50 is the combined capacity of each RAID5 set in the storage node.

For RAID50, the NSM 4150 is configured with three RAID5 sets. If the disks are 750 GB, the total capacity for that NSM 4150 equals 9 TB (12 x the single disk capacity).

RAID6

RAID6 may be thought of as RAID5 with dual parity. The dual parity of RAID6 provides fault tolerance from two drive failures in each of two RAID sets. Each array continues to operate with up to two failed drives. RAID6 significantly reduces the risk of data loss if a second hard disk drive fails while the RAID array is rebuilding.

Parity and storage capacity in RAID6

In RAID6, data is striped on a block level across a set of drives, as in RAID5, but a second set of parity is calculated and written across all the drives in the set. RAID6 provides extremely high data fault tolerance and can withstand multiple simultaneous drive failures.

A RAID6 implementation has the storage capacity of N-2 drives. A storage node configured with RAID6 has a total available storage capacity of 66% of the total system capacity. For example, in a 9 TB system, 6 TB is available for storage and 3 TB is used for overhead.




Figure 13 Parity distributed across RAID6

Drive failure and hot swapping in RAID6

The DL320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300 support RAID6 and also support hot swapping in the event of a drive failure. Hot swapping means that you can physically remove a failed drive and insert a new one without powering down the unit.

In addition to redundancy during normal operation, RAID6 further protects the RAID array against data loss during degraded mode by tolerating up to two drive failures during this vulnerable stage.

Explaining RAID devices in the RAID setup report

In the Storage category, the RAID Setup tab lists the RAID devices in the storage node and provides information about them. An example of the RAID setup report is shown in Figure 14. Information listed in the report is described in Table 9.

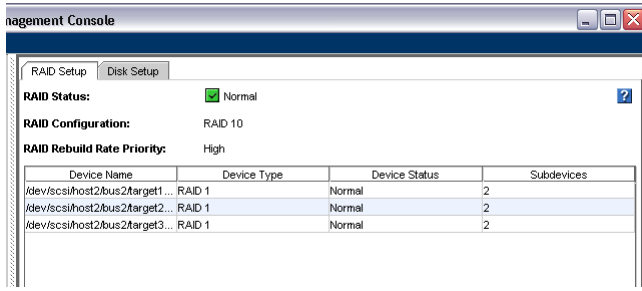


Figure 14 RAID10 in an Dell 2950

RAID devices by RAID type

Each RAID type creates different sets of RAID devices. Table 9 contains a description of the variety of RAID devices created by the different RAID types as implemented on various platform models.

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 platform and RAID level.
Device Type	The RAID level of the device.
Device Status	The RAID status of the device.
Subdevices	The number of disks included in the device. For example, in an NSM 160, RAID10 displays a Device Type of "RAID10" and subdevices as "4."

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 LeftHand Networks recommends installing VMware ESX Server on top of a server with a RAID5 or RAID6 configuration.

Devices configured in RAID0

As illustrated, if RAID0 is configured, the physical disks are combined into a single RAID disk, except for the NSM 260. In the NSM 260 with RAID0 configured, each physical disk operates as a separate RAID0 disk.




Figure 15 RAID0 an NSM 160




Figure 16 RAID0 on an NSM 260




Figure 17 RAID0 on a DL380




Figure 18 RAID0 on an IBM x3650

Devices configured in RAID10

If RAID10 is configured on storage nodes, the physical disks are combined into mirrored sets of disks and then combined into one striped disk. Examples of RAID10 devices are shown in [Figure 19](#) through [Figure 25](#). If RAID1 is configured for the NSM 260, the physical disks are combined into mirrored pairs of disks, as shown in [Figure 20](#). RAID1 uses only one pair of disks. RAID10 uses up to eight pairs of disks, depending on the platform.




Figure 19 RAID10 on an NSM 160 (mirroring done at hardware level)




Figure 20 RAID1 on an NSM 260




Figure 21 RAID10 on a DL380




Figure 22 RAID10 on the DL320s (NSM 2120) and the HP LeftHand P4500




Figure 23 RAID1+0 in the HP LeftHand P4300

	Disk	Status	Health	Safe to R
1	0	Active	normal	Yes
	1	Active	normal	Yes
	2	Active	normal	Yes
	3	Active	normal	Yes
	4	Active	normal	Yes
	5	Active	normal	Yes

1. Mirrored disk pair 1
2. Mirrored disk pair 2
3. Mirrored disk pair 3

Figure 24 Initial RAID10 setup of the Dell 2950 and NSM 2060

	Disk	Status	Health	Safe to Rem...
1	0	Active	normal	Yes
2	1	Active	normal	Yes
3	2	Active	normal	Yes
4	3	Active	normal	Yes
5	4	Active	normal	Yes
6	5	Active	normal	Yes
7	6	Active	normal	Yes
8	7	Active	normal	Yes
	8	Active	normal	Yes
	9	Active	normal	Yes
	10	Active	normal	Yes
	11	Active	normal	Yes
	12	Active	normal	Yes
	13	Active	normal	Yes
8	14	Hot spare	normal	Yes

1. Mirrored disk pair 1
2. Mirrored disk pair 2
3. Mirrored disk pair 3
4. Mirrored disk pair 4
5. Mirrored disk pair 5
6. Mirrored disk pair 6
7. Mirrored disk pair 7
8. Hot spare

Figure 25 Initial RAID10 setup of the NSM 4150

NOTE:

The initial disk setup shown above for the NSM 4150 may change over time if you have to replace hot-swap disks.

Devices configured in RAID5

If RAID5 is configured, the physical disks are grouped into one or more RAID5 sets.




Figure 26 RAID5 set in an NSM 160




Figure 27 RAID5 + spare in an NSM 160




Figure 28 RAID5 set in a DL380




Figure 29 RAID5 set in the DL320s (NSM 2120), and the HP LeftHand P4500




Figure 30 RAID5 set in the HP LeftHand P4300





Figure 31 RAID5 set in a IBM x3650



Disk	Status	Health	Safe to R
0	Active	normal	Yes
1	Active	normal	Yes
2	Active	normal	Yes
3	Active	normal	Yes
4	Active	normal	Yes
5	Active	normal	Yes

1. RAID5 set

Figure 32 Initial RAID5 setup of the Dell 2950 and NSM 2060

 **NOTE:**

The initial disk setup shown above for the Dell 2950 and NSM 2060 may change over time if you have to replace disks.

RAID5 in the NSM 260

RAID5 in the NSM 260 consists of either six-disk sets, shown in [Figure 33](#), or sets of five disks plus a spare so that the single disk acts as a hot spare for the RAID set, shown in [Figure 34](#).

RAID50 in the NSM 4150 consists of three RAID5 sets using all 15 disks, shown in [Figure 35](#).




Figure 33 NSM 260 RAID5 using six-disk sets




Figure 34 NSM 260 RAID5 using five disks plus a hot spare

	Disk	Status	Health	Safe to ...
1	0	Active	normal	Yes
	1	Active	normal	Yes
	2	Active	normal	Yes
	3	Active	normal	Yes
	4	Active	normal	Yes
2	5	Active	normal	Yes
	6	Active	normal	Yes
	7	Active	normal	Yes
	8	Active	normal	Yes
	9	Active	normal	Yes
3	10	Active	normal	Yes
	11	Active	normal	Yes
	12	Active	normal	Yes
	13	Active	normal	Yes
	14	Active	normal	Yes

1. RAID5 set
2. RAID5 set
3. RAID5 set

Figure 35 Initial RAID50 setup of the NSM 4150



NOTE:

The initial disk setup shown above for the NSM 4150 may change over time if you have to replace disks.

RAID6 in the DL320s (NSM 2120), HP LeftHand P4500

For RAID6, the physical disks are grouped into sets. RAID6 uses two sets of disks.




Figure 36 DL320s (NSM 2120) and HP LeftHand P4500 RAID6 using two six-disk sets

RAID6 in the HP LeftHand P4300

RAID6 in the HP LeftHand P4300 is striped with parity into a single array.




Figure 37 Raid5 in P4300

Planning the RAID configuration

The RAID configuration you choose for the storage node depends on your plans for data fault tolerance, data availability, and capacity growth. If you plan to expand your network of storage nodes and create clusters, choose your RAID configuration carefully.

△ CAUTION:

After you have configured RAID, you cannot change the RAID configuration without deleting all data on the storage node.

Data replication

Keeping multiple copies of your data can ensure that data will be safe and will remain available in the case of disk failure. There are two ways to achieve data replication:

- Configure RAID1, 10, 5, 5 + spare, 50, or 6 within each storage node to ensure data redundancy.
- Always replicate data volumes across clusters of storage nodes, regardless of RAID level, for added data protection and high availability.

Using RAID for data redundancy

Within each storage node, RAID1 or RAID10 can ensure that two copies of all data exist. If one of the disks in a RAID pair goes down, data reads and writes can continue on the other disk. Similarly, RAID5, RAID50, or RAID6 provides redundancy by spreading parity evenly across the disks in the set.

If one disk in a RAID5 or RAID6 set goes down, data reads and writes continue on the remaining disks in the set. In RAID50, up to one disk in each RAID5 set can go down, and data reads and writes continue on the remaining disks.

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

📝 NOTE:

If you plan to create all data volumes on a single storage node, use RAID1/10, RAID5 or RAID6 to ensure data redundancy within that storage node.

Using volume replication in a cluster

A cluster is a group of storage nodes across which data can be replicated. Volume replication across a cluster of storage nodes protects against disk failures within a storage node, failure of an entire storage node or external failures like networking or power. For example, if a single disk or an entire storage node in a cluster goes offline, data reads and writes can continue because an identical copy of the volume exists on other storage nodes in the cluster.

Using RAID with replication in a cluster

Always use replication in a cluster to replicate volumes across storage nodes. The redundancy provided by RAID10, 5, 50, or 6 ensures availability at the storage node level. Replication of volumes in a cluster ensures availability at the cluster level. For example,

- Using replication, up to three copies of a volume can be created on a cluster of three storage nodes. The replicated configuration ensures that two of the three storage nodes can go offline and the volume will still be accessible.
- Configuring RAID10 on these storage nodes means that each of these three copies of the volume is stored on two disks within the storage node, for a total of six copies of each volume. For a 50 GB volume, 300 GB of disk capacity is used.

RAID5/50 uses less disk capacity than RAID 1 / 10, so it can be combined with replication and still use capacity efficiently. One benefit of configuring RAID5/50 in storage nodes that use replication in a cluster is that if a single disk goes down, the data on that storage node can be rebuilt using RAID instead of requiring a complete copy from another storage node in the cluster. Rebuilding the disks within a single set is faster and creates less of a performance hit to applications accessing data than copying data from another storage node in the cluster.

RAID6 provides similar space benefits of RAID5 with the additional protection of being able to survive the loss of up to two drives.



NOTE:

If you are replicating volumes across a cluster: Configuring the storage node for RAID1/10 consumes half the capacity of the storage node. Configuring the storage node for RAID5/50 provides redundancy within each storage node while allowing most of the disk capacity to be used for data storage. RAID6 provides greater redundancy on a single storage node, but consumes more disk space than RAID5.

Table 10 summarizes the differences in data availability and safety of the different RAID levels on stand-alone storage nodes versus on those RAID levels with replicated volumes in a cluster.

Table 10 Data availability and safety in RAID configurations

Configuration	Safety and availability during disk failure	Data availability if entire storage node fails or if network connection to storage node lost
Stand-alone storage nodes, RAID0	No	No
Stand-alone storage nodes, RAID1 / 10, RAID10 + spare	Yes. In any configuration, 1 disk per mirrored pair can fail.	No

Configuration	Safety and availability during disk failure	Data availability if entire storage node fails or if network connection to storage node lost
Stand-alone storage nodes, RAID5, 5 + spare, 50	Yes, for 1 disk per array	No
Stand-alone storage nodes, RAID6	Yes, for 2 disks per array	No
Replicated volumes on clustered storage nodes, RAID 0	Yes. However, if any disk in the storage node fails, the entire storage node must be copied from another storage node in the cluster.	Yes
Replicated volumes on clustered storage nodes, RAID 5 / 50	Yes. 1 disk per RAID set can fail without copying from another storage node in the cluster.	Yes
Replicated volumes on clustered storage nodes, RAID6	Yes. 2 disks per RAID set can fail without copying from another storage node in the cluster.	Yes
Replicated volumes on clustered VSAs with virtual RAID	Depends on the underlying RAID configuration of the platform on which the VSA is installed. HP recommends configuring RAID5 or RAID6.	Yes, if underlying platform configured for RAID other than RAID0.

Mixing RAID configurations

You may mix storage nodes with different configurations of RAID within a cluster. This allows you to add new storage nodes with different RAID levels. Be certain to calculate the capacity of additional storage nodes running the desired RAID level, because the cluster operates at the smallest usable per-storage node capacity.

For instance, you have four 1 TB NSM 160s running RAID10. You purchase two additional 1 TB NSM 160s which you want to run with RAID5.

In your existing cluster, a single 1 TB NSM 160 in RAID10 provides 0.5 TB usable storage. A single NSM 160 in RAID5 provides 0.75 TB usable storage. However, due to the restrictions of how the cluster uses capacity, the NSM 160 in RAID5 will still be limited to 0.5 TB per storage node.

Your best choice in this situation might be to configure the NSM 160 as RAID5 + a spare, thus using most of the available storage and enhancing the reliability of the cluster.

Setting RAID rebuild rate

Choose the rate at which the RAID configuration rebuilds if a disk is replaced. The RAID rebuild rate is set as a priority against other operating system tasks, except for the NSM 260 platform, for which the rebuild rate is a percentage of the throughput of the RAID card.

NOTE:

You cannot set the RAID rebuild rate on a VSA, since there is nothing to rebuild.

General Guidelines

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

△ CAUTION:

In the IBM x3650, the RAID rebuild rate cannot currently be changed from high. This setting may affect the SAN if RAID10 or RAID5 needs to be rebuilt.

Set RAID rebuild rate

1. In the navigation window, log in to a storage node and select the Storage category.
2. On the RAID Setup tab, click RAID Setup Tasks and select the RAID Rebuild Rate Priority choice.
The RAID Rebuild Rate Priority window opens. This window will be different for different platforms, as described above.
3. Change the rebuild settings as desired.
4. Click OK.

The settings are then ready when and if RAID rebuild takes place.

Reconfiguring RAID

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

Requirements for reconfiguring RAID

Changing preconfigured RAID on a new storage node

RAID must be reconfigured on individual storage nodes before they are added to a management group. If you want to change the preconfigured RAID level of a storage node, you must make the change before you add the storage node to a management group.

Changing RAID on storage nodes in management groups

You cannot reconfigure RAID on a storage node that is already in a management group. If you want to change the RAID configuration for a storage node 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, expand the configuration categories for the storage node 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 node may take several hours for the disks to synchronize in a RAID10, or RAID5/50, or RAID6 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. The storage node is ready for data transfer at this point.

Monitoring RAID status

RAID is critical to the operation of the storage node. If RAID has not been configured, the storage node cannot be used. Monitor the RAID status of a storage node to ensure that it remains normal. If the RAID status changes, a CMC alert is generated. You can configure additional alerts to go to an email address or to an SNMP trap. See [“Using alerts for active monitoring”](#) on page 135 for instructions to set these additional alerts.

Data transfer and RAID status

A RAID status of Normal, Rebuild, or Degraded all allow data transfer. The only time data cannot be transferred to or from the storage node is if the RAID status shows Off.

Data redundancy and RAID status

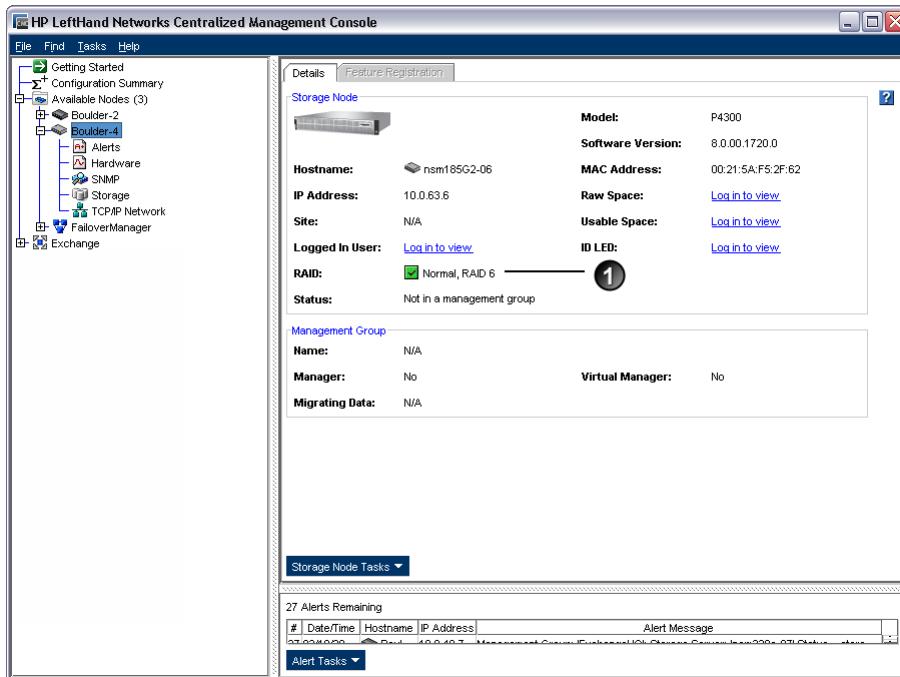
In a RAID1/10 or RAID5/50 configuration, when RAID is degraded, there is no full data redundancy. Therefore, data is at risk if there is a disk failure when RAID is degraded.

In RAID6, 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 RAID1 / 10 configuration, loss of a second disk within a pair will result in data loss. In a degraded RAID5 configuration, loss of a second disk will result in data loss. In a degraded RAID50 configuration, loss of a second disk in a single RAID5 set will result in data loss. In a degraded RAID6 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 displays in the Details Tab on the main CMC window when a storage node is selected in the navigation window.



1. RAID status

Figure 38 Monitoring RAID status on the main CMC window

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 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 platforms (NSM 160, NSM 260, DL380, DL320s [NSM 2120], Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300), simply replace the faulty, inactive, uninitialized, or missing disk.
For non-hot-swap platforms (IBM x3650) you must add a disk to RAID using Storage > Disk Setup tab if you are inserting a replacement disk.
- **Off**—Data cannot be stored on the storage node. The storage node 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.

△ CAUTION:

The IBM x3650 does not support hot-swapping disk drives.
Hot-swapping drives is NOT supported for RAID 0 on any platform.

Table 11 Disk management tasks for storage nodes

Disk setup function	Available in model
Monitor disk information	All
Power on or off a disk	<ul style="list-style-type: none">IBM x3650
Add a disk to RAID	<ul style="list-style-type: none">NSM 260 (use only for adding capacity/arrays to a chassis)

Getting there

1. In the navigation window, select a storage node.
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 node.

Figure 39 shows the Disk Setup tab and Table 12 describes the corresponding disk report.

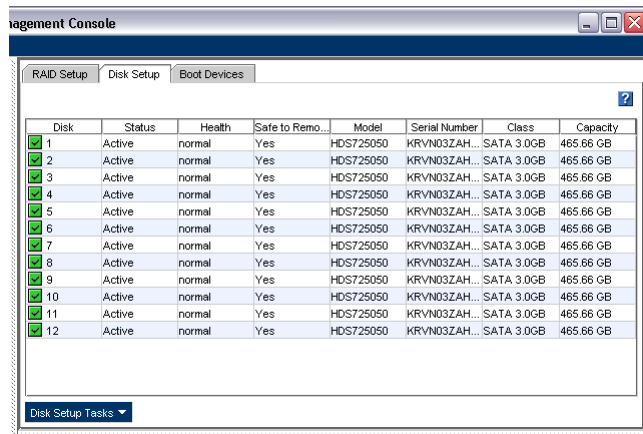


Figure 39 Example of columns in the Disk Setup tab

Table 12 Description of items on the disk report

This item	Describes this
Disk	Corresponds to the physical slot in the storage node.

This item	Describes this
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 possible”) • Faulty (predictive failure status indicating “don’t wait to replace”)
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 you are preparing to replace them.

Viewing disk status for the NSM 160

The disks are labeled 1 through 4 in the Disk Setup window, [Figure 40](#), and correspond with the disk drives from left to right (1 through 4) as shown in [Figure 41](#).

For the NSM 160, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

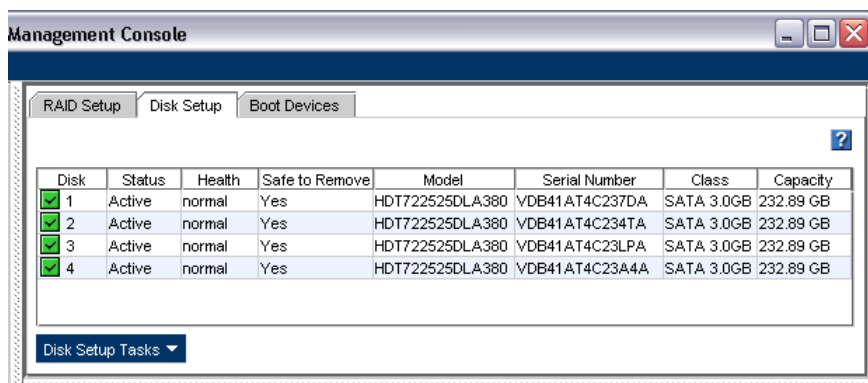


Figure 40 Viewing the Disk Setup tab in an NSM 160



Figure 41 Diagram of the drive bays in the NSM 160

Viewing disk status for the NSM 260

For the NSM 260, the disks are labeled 1 through 12 in the Disk Setup window and correspond with the disk drives from left to right (1, 2, 3, 4 top row: 5, 6, 7, 8 middle row: 9, 10, 11, 12 bottom row) and top to bottom when you are looking at the front of the NSM 260.

For the NSM 260, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

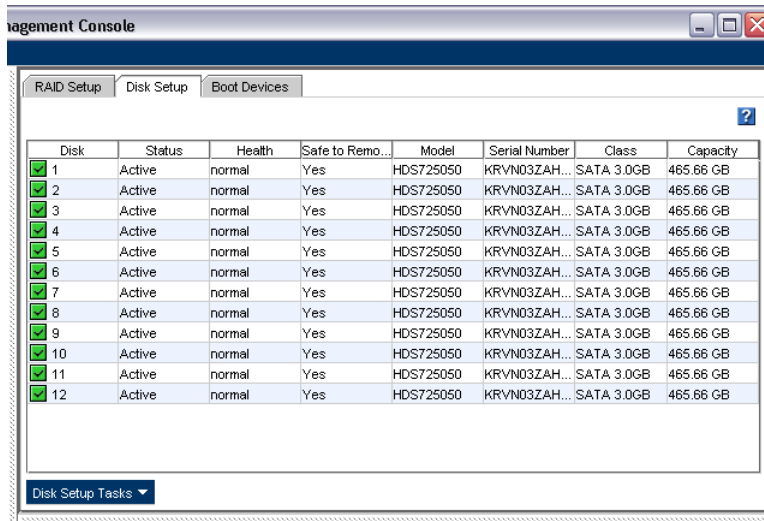


Figure 42 Viewing the Disk Setup tab in an NSM 260




Figure 43 Diagram of the drive bays in the NSM 260

Viewing disk status for the DL380

For the DL380, the disks are labeled 0 through 5 in the Disk Setup window, shown in Figure 44, and correspond with the disk drives from left to right (5-3-1 on the top and 4-2-0 on the bottom), as shown in Figure 45 when you are looking at the front of the DL380.

For the DL380, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

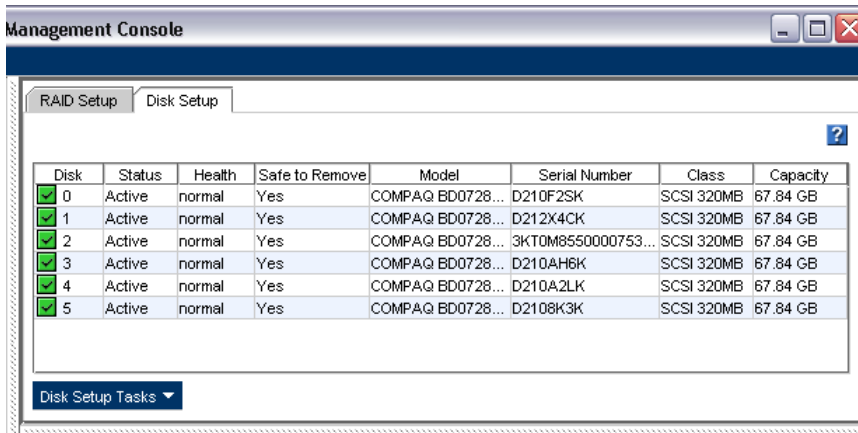


Figure 44 Viewing the Disk Setup tab in a DL380



Figure 45 Arrangement of drives in the DL380

Viewing disk status for the DL320s (NSM 2120)

The disks are labeled 1 through 12 in the Disk Setup window, shown in Figure 46, and correspond with 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 47 when you are looking at the front of the DL320s (NSM 2120).

For the DL320s (NSM 2120), the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

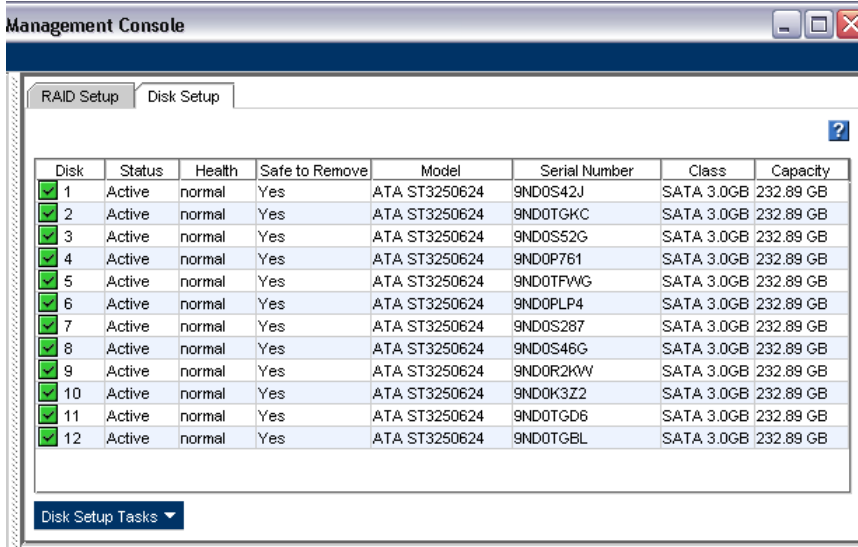


Figure 46 Viewing the Disk Setup tab in a DL320s (NSM 2120)



Figure 47 Diagram of the drive bays in a DL320s (NSM 2120)

Viewing disk status for the IBM x3650

For the IBM x3650, the disks are labeled 0 through 5 in the Disk Setup window, [Figure 48](#), and correspond with the disk drives labeled 0 through 5 from left to right, top to bottom, when you are looking at the front of the IBM x3650, [Figure 49](#).

For the IBM x3650, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

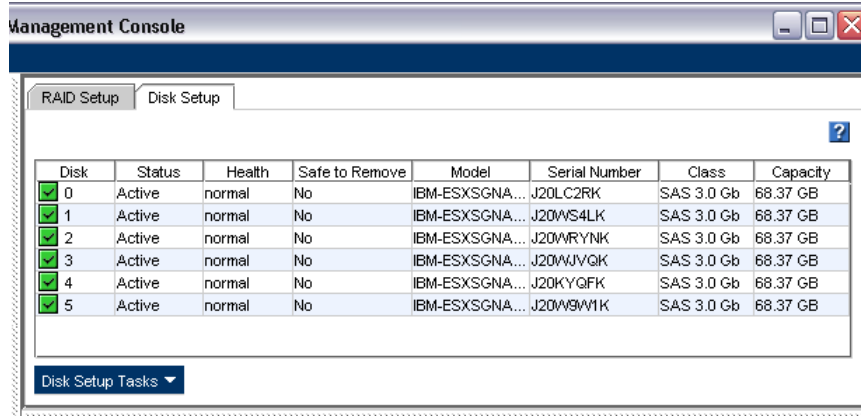


Figure 48 Viewing the Disk Setup tab in the IBM x3650



Figure 49 Arrangement of drives in the IBM x3650

Viewing disk status for the VSA

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




Figure 50 Viewing the disk status of a VSA

NOTE:

If you want to change the size of the data disk in a VSA, see the *HP LeftHand P4000 VSA User Manual* for instructions about recreating the disk in the VI Client.

Viewing disk status for the Dell 2950 and NSM 2060

The disks are labeled 0 through 5 in the Disk Setup window and correspond with the disk drives from left to right (0-2-4 on the top row, and 1-3-5 on the bottom row) when you are looking at the front of the Dell 2950 or NSM 2060.

For the Dell 2950 or NSM 2060, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

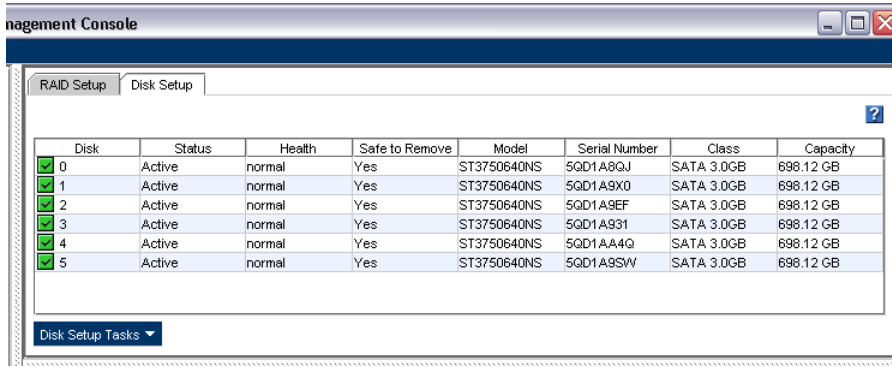


Figure 51 Viewing the Disk Setup tab in a Dell 2950 or NSM 2060



Figure 52 Drive bays, with bezel on, in a Dell 2950 or NSM 2060



Figure 53 Drive bays, with bezel off, in a Dell 2950 or NSM 2060

Viewing disk status for the NSM 4150

The disks are labeled 0 through 14 in the Disk Setup window and correspond with the disk drives from left to right when you are looking at the front of the NSM 4150.

For the NSM 4150, the columns Health and Safe to Remove will respectively help you assess the health of a disk, and whether or not you can replace it without losing data.

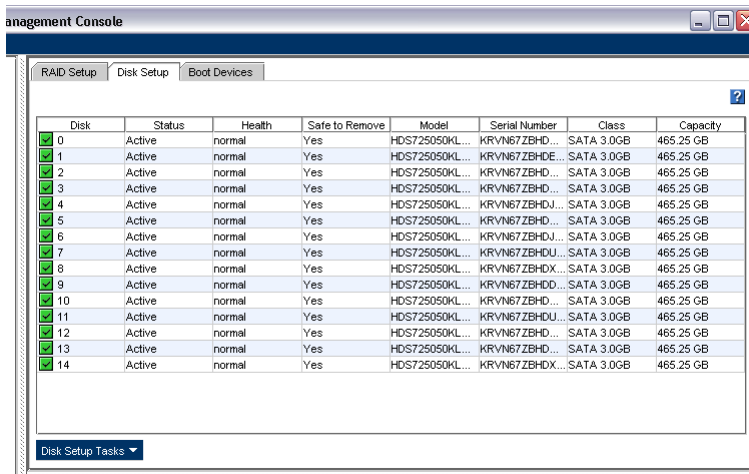


Figure 54 Viewing the Disk Setup tab in a NSM 4150




Figure 55 Drive bays, with bezel on, in an NSM 4150




Figure 56 Drive bays, with bezel off, in an NSM 4150

Viewing disk status for the HP LeftHand P4500

The disks are labeled 1 through 12 in the Disk Setup window [Figure 57](#) and correspond with 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 58](#) when you are looking at the front of the HP LeftHand P4500.

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

Disk	Status	Health	Safe to Remo...	Model	Serial Number	Class	Capacity
1	Active	normal	Yes	HP DF0450B...	JMV6703C	SAS 3.0GB	419.19 GB
2	Active	normal	Yes	HP DF0450B...	JMV2ELBC	SAS 3.0GB	419.19 GB
3	Active	normal	Yes	HP DF0450B...	JMV3EJMC	SAS 3.0GB	419.19 GB
4	Active	normal	Yes	HP DF0450B...	JMV2Z9VC	SAS 3.0GB	419.19 GB
5	Active	normal	Yes	HP DF0450B...	JMV46S7C	SAS 3.0GB	419.19 GB
6	Active	normal	Yes	HP DF0450B...	JMV2EV4C	SAS 3.0GB	419.19 GB
7	Active	normal	Yes	HP DF0450B...	JMV3H60C	SAS 3.0GB	419.19 GB
8	Active	normal	Yes	HP DF0450B...	JMV3HBPC	SAS 3.0GB	419.19 GB
9	Active	normal	Yes	HP DF0450B...	JMV3EL2C	SAS 3.0GB	419.19 GB
10	Active	normal	Yes	HP DF0450B...	JMV65N2C	SAS 3.0GB	419.19 GB
11	Active	normal	Yes	HP DF0450B...	JMV3EVTTC	SAS 3.0GB	419.19 GB
12	Active	normal	Yes	HP DF0450B...	JMV3YEMC	SAS 3.0GB	419.19 GB

Figure 57 Viewing the Disk Setup tab in a HP LeftHand P4500

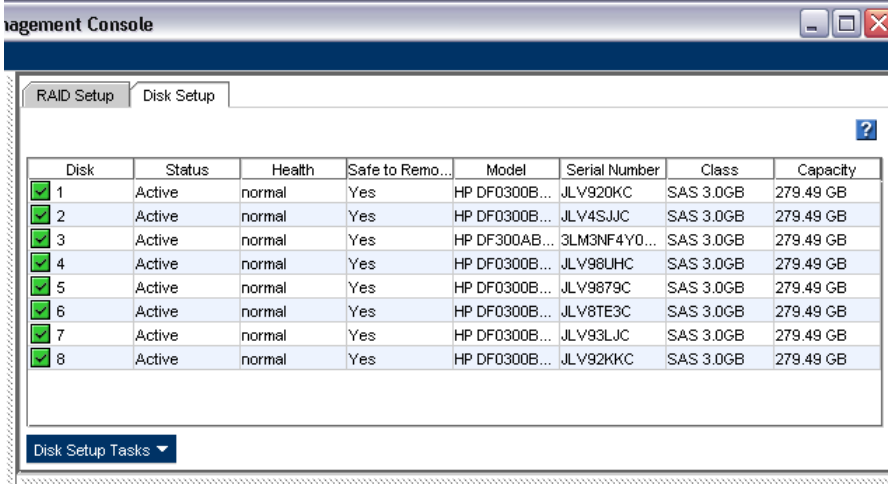


Figure 58 Diagram of the drive bays in a HP LeftHand P4500

Viewing disk status for the HP LeftHand P4300

The disks are labeled 1 through 8 in the Disk Setup window [Figure 59](#) and correspond with 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) shown in [Figure 60](#) when you are looking at the front of the HP LeftHand P4300.

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



The screenshot shows the Management Console interface with the 'Disk Setup' tab selected. A table lists the status of 8 disks. All disks are active and in normal health, with a capacity of 279.49 GB each.

Disk	Status	Health	Safe to Remo...	Model	Serial Number	Class	Capacity
1	Active	normal	Yes	HP DF0300B...	JLV920KC	SAS 3.0GB	279.49 GB
2	Active	normal	Yes	HP DF0300B...	JLV4SJJC	SAS 3.0GB	279.49 GB
3	Active	normal	Yes	HP DF300AB...	3LM3NF4Y0...	SAS 3.0GB	279.49 GB
4	Active	normal	Yes	HP DF0300B...	JLV98UHC	SAS 3.0GB	279.49 GB
5	Active	normal	Yes	HP DF0300B...	JLV9879C	SAS 3.0GB	279.49 GB
6	Active	normal	Yes	HP DF0300B...	JLV8TE3C	SAS 3.0GB	279.49 GB
7	Active	normal	Yes	HP DF0300B...	JLV93LJC	SAS 3.0GB	279.49 GB
8	Active	normal	Yes	HP DF0300B...	JLV92KKC	SAS 3.0GB	279.49 GB

Figure 59 Viewing the Disk Setup tab in a HP LeftHand P4300



Figure 60 Diagram of the drive bays in a HP LeftHand P4300

Replacing a disk

The procedures for replacing a disk are different for various platforms:

- RAID0 in all platforms
- Platforms that support hot-swapping drives, which include the NSM 160, NSM 260, DL380, and DL320s (NSM 2120), Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300
- Non-hot-swap platforms, including the IBM x3650

In the VSA

If you are replacing a disk on a server that is hosting the VSA, refer to the manufacturer's instructions.

If you want to change the disk size on a VSA, you must recreate the hard disk in the VI Client. See the *VSA User Manual* for detailed instructions.

Using Repair Storage Node

In certain circumstances, you may have to use the Repair Storage Node feature when replacing disks. Such circumstances include the following:

- RAID is OFF on a storage node with RAID0

- When replacing multiple disks on a storage node with RAID5/50 or RAID6
- When multiple disks on the same mirror set need to be replaced on a storage node with RAID10.

See “[Replacing disks appendix](#)” on page 327 for further information.

Overview of replacing a disk

The correct procedure for replacing a disk in a storage node depends upon a number of factors, including the RAID configuration, the replication level of volumes and snapshots, and the number of disks you are replacing. Unless you are replacing a disk in a storage node that is not in a cluster, data must be rebuilt either just on the replaced disk or, in the case of RAID0, on the entire storage node.

Replacing a disk in a storage node includes the following steps:

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

See [Using Repair Storage Node](#) on page 80 for situations in which you can use this feature to save unnecessary restripes of your data.

Replacing disks in hot-swap platforms

In hot-swap platforms running RAID1/10, 5, 50, or 6, you can remove a faulty or failed disk and replace it with a new one. RAID will rebuild and the drive will return to Normal status.

△ CAUTION:

Before replacing a drive in a hot-swap platform, 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 RAID1/10, RAID5, RAID50, or RAID6, all drives indicate they are safe to remove. However, you should only hot-swap one drive at a time. If it is necessary to replace more than one drive, always check the Safe to Remove status again. You must 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 you may replace it.

For example, if an array is Rebuilding, no other drives in the array (except for hot-spare drives) will be safe to remove. However, if the configuration includes two or more arrays and those arrays are Normal, the Safe To Remove status will indicate that drives in those other arrays may be replaced. Note that the Safe To Remove status will always be No when in a RAID0 configuration until the drive is powered off. Also note that Hot Spare, Inactive, and Uninitialized drives are always safe to remove.

Replacing disks in non-hot-swap platforms (IBM x3650)

In non-hot-swap platforms running RAID1/10 or 5, you must power off in the CMC the faulty or failed disk before you physically replace the disk in the chassis. After physically replacing the disk you must power on in the CMC the newly replaced disk.

Replacing disks in RAID0 configurations

In RAID0 configurations, you must power off in the CMC the faulty or failed disk before you physically replace the disk in the chassis. After physically replacing the disk you must power on in the CMC the newly replaced disk. See “[Best practice checklist for single disk replacement in RAID0](#)” on page 82.

△ CAUTION:

If you remove a disk from a RAID0 configuration, all the data on the storage node will be lost. The Best Practice Checklist describes how to prepare for a single disk replacement in RAID0.

Preparing for a disk replacement

Use this section if you are replacing 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](#)” on page 327 for these situations:

- If RAID has gone 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

How you prepare for a disk replacement differs according to the RAID level of the storage node and whether it is a hot-swap platform. You should carefully plan any disk replacement to ensure data safety, regardless of whether the platform is hot-swap. The following checklists outline steps to help ensure your data remains safe while you replace a disk.

Identify physical location of storage node and disk

Before you begin the disk replacement process, you should identify the physical location of both the storage node in the rack and the disk in the storage node.

- Know the name and physical location of the storage node that needs the disk replacement.
- Know the physical position of the disk in the storage node. See “[Verifying disk status](#)” on page 74 for diagrams of disk layout in the various platforms.
- 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 RAID0

△ CAUTION:

Do not use hot-swap procedures on any storage node running in RAID0.

In RAID0, always power off the drive in the CMC before removing it. RAID0 provides no fault tolerance by itself, so when you do power off the drive you lose the data on the storage node. Therefore, if you need to replace a disk in a RAID0 configuration, HP recommends the following:

- All volumes and snapshots have a minimum of 2-way replication.
- If volumes or snapshots are not replicated, change them to 2-way replication 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, 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](#)” on page 327 if you have more than one disk to replace, or if you are unsure which disk needs replacing.

Best practice checklist for single disk replacement in RAID1/10, RAID5, RAID50, and RAID6

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

- All volumes and snapshots should show a status of Normal.
- Any volumes or snapshots that were being deleted have completed deletion.
- RAID status is Normal, or
- If RAID is Rebuilding or Degraded, for platforms that support hot-swapping of drives, the Safe to Remove column indicates “Yes” that the drive can safely be replaced.

Replacing a disk in RAID0

Complete the checklist for single disk replacement RAID0.

Manually power off the disk in the CMC for RAID0

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

1. In the navigation window, select the storage node in which you want to replace the disk.
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 node

See the hardware documentation for the storage node.

Manually powering on the disk in the CMC

When you must insert the new disk into the storage node, 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, like this —. [Figure 61](#) shows a representation of a missing disk in a storage node.

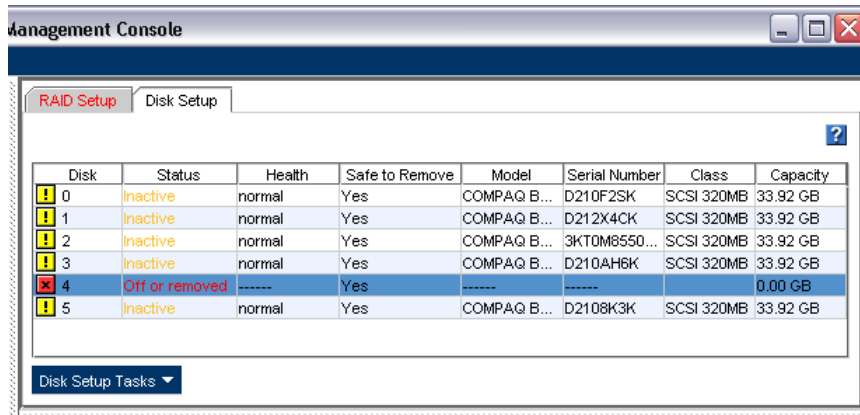


Figure 61 Viewing a power off or missing disk

1. In the navigation window, select the storage node in which you replaced the disk drive.
2. Select the Storage category in the 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 goes to Normal. Volumes start restriping on the entire storage node. 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 non-hot-swap platform (IBM x3650)

Complete the checklist for the RAID1/10 or RAID5 level disk replacement. Then follow the procedures below.

△ CAUTION:

IBM x3650: You must always use a new drive when replacing a disk in an IBM x3650. Never reinsert the same drive and power it on again.

In non-hot-swap platforms running RAID1/10 or 5, you must power off in the CMC the faulty or failed disk before you physically replace the disk in the chassis. After physically replacing the disk, power on in the CMC the newly replaced disk.

Manually power off the disk in the CMC for RAID1/10 and RAID5

You first power off in the CMC the disk you are replacing. Powering off a single disk in RAID1/10, or RAID5 causes RAID to run in a degraded state.

1. In the navigation window, select the storage node in which you want to replace the disk.

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 node

See the hardware documentation for the storage node.

Manually powering on the disk in the CMC

When you must insert the new disk into the storage node, 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, like this —. [Figure 62](#) shows a representation of a missing disk in a storage node.

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

Figure 62 Viewing a power off or missing disk

Manually powering on the disk in the IBM x3650

1. In the navigation window, select the IBM x3650 in which you replaced the disk drive.
2. Select the Storage configuration category.

△ CAUTION:

Wait until the RAID status on the RAID Setup tab displays “Rebuilding.”

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.

RAID rebuilding

After the disk is powered on, RAID starts rebuilding on the replaced disk. Note that 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.

Replacing a disk in a hot-swap platform (NSM 160, NSM 260, DL380, DL320s [NSM 2120], Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300)

Complete the checklist for replacing a disk in RAID1/10, 5, 50, or RAID6. Then follow the appropriate procedures for the platform.

△ CAUTION:

You must always use a new drive when replacing a disk in an Dell 2950, NSM 2060, or NSM 4150. Never reinsert the same drive or another drive from the same Dell 2950, NSM 2060, or NSM 4150.

Replace the disk

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

Physically replace the disk drive in the storage node

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

RAID rebuilding

After the disk is replaced, RAID starts rebuilding on the replaced disk. Note that 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.

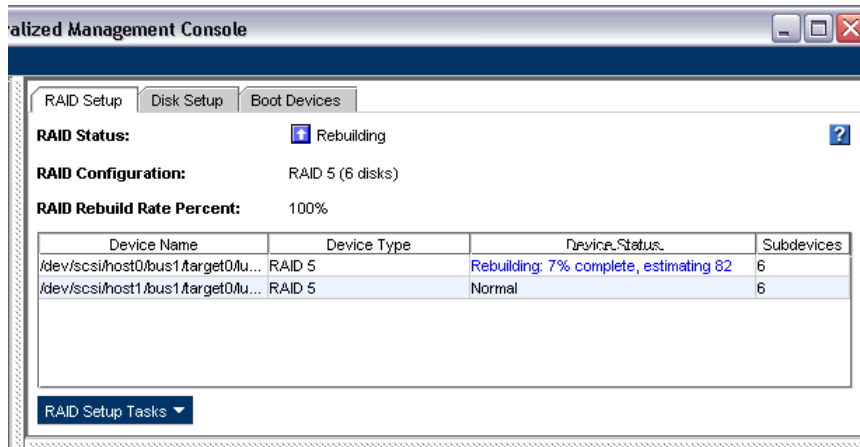


Figure 63 RAID rebuilding on the RAID Setup tab

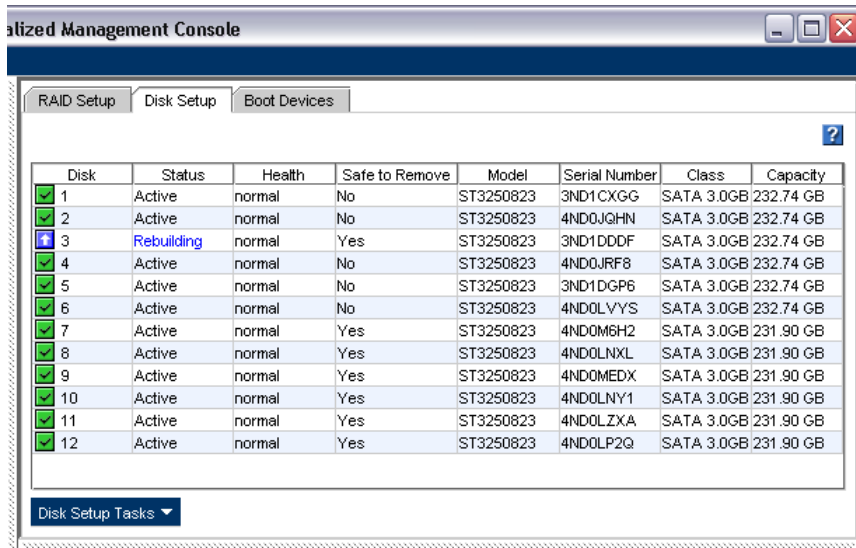


Figure 64 Disk rebuilding on the Disk Setup tab

4 Managing the network

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

- Configure the individual TCP/IP interfaces.
- Set up and manage a DNS server.
- Manage the routing table.
- View and configure the TCP interface speed and duplex, frame size and NIC flow control.
- Update the list of managers running in the management group to which a storage node belongs.
- Bond NICs to ensure continuous network access or to improve bandwidth.

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 node before creating a management group, or before adding the storage node to a management group and cluster.
- Use static IP addresses, or reserved addresses if using DHCP.
- Configure storage node settings for speed and duplex, frame size, and flow control BEFORE bonding NICs and before putting the storage node into a management group and cluster.
- If adding a second IP address to a storage node, the second IP address must be on a separate subnet. If the two IP addresses are on the same subnet, they must be bonded.

Changing network configurations

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

If you do need to change the network characteristics of a storage node while it is in a cluster, be sure to follow our recommended best practices.

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 node at a time.

- Some network changes cause the storage server to restart the SAN/iQ services, making the storage node unavailable for a short time. Check the Availability tab for each storage node to see if any volumes will become unavailable if the services restart on the storage node.
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 node.
2. Open the tree under the storage node and select TCP/IP Network.

The TCP/IP tab


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

Identifying the network interfaces

A storage node 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 node.

Table 13 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 storage node's the serial port, as described in Chapter 23 on page 341.

Table 13 Identifying the network interfaces on the storage node

Ethernet interfaces	
Where labeled	Labeled as one of the these
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 node's serial port	<ul style="list-style-type: none"> • Intel Gigabit Ethernet • Broadcom Gigabit Ethernet
On the label on the back of the storage node	<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 node 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 node and open the tree below it.
2. Select the TCP/IP Network category.
3. Select 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 you want 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 node 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 you want 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 node.

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 nodes on the DHCP server. This is required because management groups use unicast communication.

To set IP address using DHCP

1. Select from the list the interface you want 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

Network interface bonding provides high availability, fault tolerance, load balancing and/or bandwidth aggregation for the network interface cards in the storage node. Bonds are created by “bonding” 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 node 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/or 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 node into separate switches. While Link Aggregation Dynamic Mode will only survive a port failure, Active-Passive will survive a switch failure.

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

- These settings are controlled on the TCP Status tab 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 node is set for Auto/Auto, the switch must be set the same. See “[The TCP status tab](#)” on page 107 for more information.

Table 14 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

Feature	Active-passive	Link aggregation dynamic mode	Adaptive load balancing
Supported storage nodes	NSM160	NSM160	NSM160
	NSM260	NSM260	NSM260
	DL 380	DL 380	DL 380
	DL 320s (NSM 2120)	DL 320s (NSM 2120)	DL 320s (NSM 2120)
	IBM x3650	IBM x3650	IBM x3650
	Dell 2950	Dell 2950	Dell 2950
	NSM 2060	NSM 2060	NSM 2060
	NSM 4150	NSM 4150	NSM 4150
	HP LeftHand P4500	HP LeftHand P4500	HP LeftHand P4500
HP LeftHand P4300	HP LeftHand P4300	HP LeftHand P4300	

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

How active-passive 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 node are labeled as listed in [Table 15](#). 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 15 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
Slot1:Port0 [NSM 260]	Physical interface in a PCI slot. This interface acts as a 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 16 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 17 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 18 shows the states of Eth0 and Eth1 when configured for Active-Passive are shown below.

Table 18 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

Failover status	Status of Eth0	Status of Eth1
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 configurations with active-passive

Two simple network configurations using Active-Passive in high availability environments are illustrated.

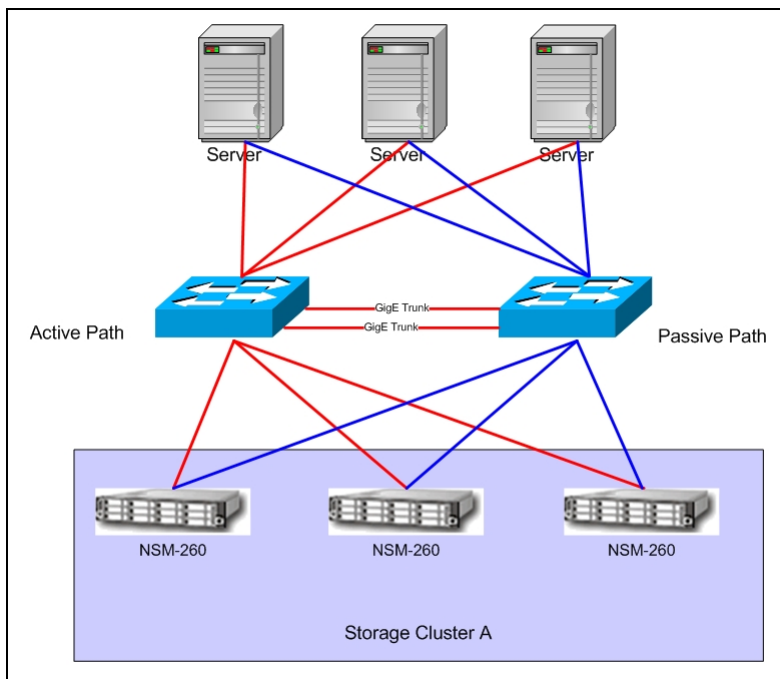


Figure 65 Active-passive in a two-switch topology with server failover

The two-switch scenario in [Figure 65](#) is a basic, yet effective, method for ensuring high availability. If either switch failed, or a cable or NIC on one of the storage nodes failed, the Active-Passive bond would cause the secondary connection to become active and take over.




Figure 66 Active-passive failover in a four-switch topology

[Figure 66](#) illustrates the Active-Passive configuration in a four-switch topology.

How link aggregation dynamic mode works

Link Aggregation Dynamic Mode allows the storage node 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 node 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 are 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 19 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 20 shows the states of Eth0 and Eth1 when configured for Link Aggregation Dynamic Mode.

Table 20 NIC status during failover with link aggregation dynamic mode

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

Example network configurations with link aggregation dynamic mode

A simple network configuration using Link Aggregation Dynamic Mode in a high availability environment is illustrated.




Figure 67 Link aggregation dynamic mode in a single-switch topology

How adaptive load balancing works

Adaptive Load Balancing allows the storage node 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 are 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 21 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 22 shows the states of Motherboard:Port1 and Motherboard:Port2 when configured for Adaptive Load Balancing.

Table 22 NIC status during failover with Adaptive Load Balancing

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

Example network configurations with adaptive load balancing

A simple network configuration using Adaptive Load Balancing in a high availability environment is illustrated.




Figure 68 Adaptive Load Balancing in a two-switch topology

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”](#) on page 108.

Bond guidelines

- Create a bond on a storage node before you add the storage node to a management group.
- Create bonds of 2 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 node before you add it to a management group.
- Verify that the bond is created.

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

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

See “[Deleting a NIC bond](#)” on page 344 for information about deleting NIC bonds using the Configuration Interface.

Creating the bond

1. Log in to the storage node.
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.
7. Enter the Subnet mask.
8. (Optional) Enter the default gateway.
9. Click OK.

 **NOTE:**

The storage node 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 node.

10. Click OK to confirm the TCP/IP changes.

A message opens, prompting you to search for the bonded storage node on the network.

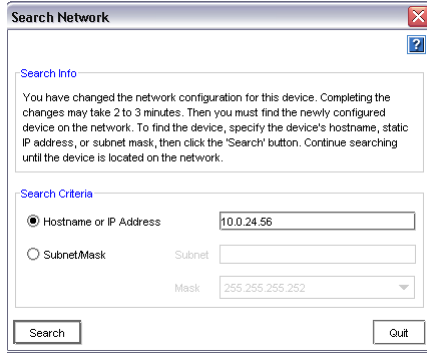


Figure 69 Searching for the bonded storage node on the network

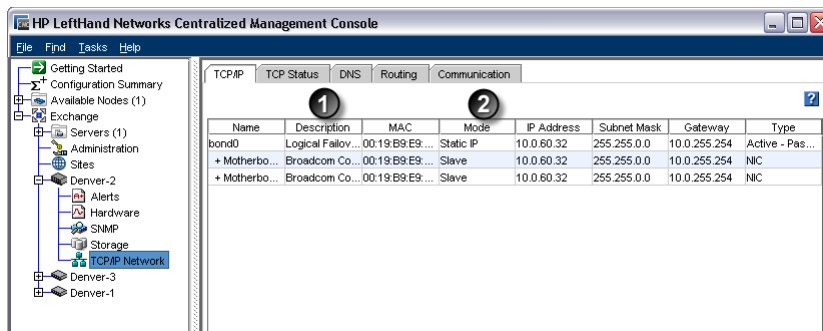
11. Search for the storage node by Host Name or IP address, or by Subnet/mask.



NOTE:

Because it can take a few minutes for the storage node 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.



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

Figure 70 Viewing a new active-passive bond

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 node and open the tree below it.

2. Select the TCP/IP Network category and click Communication tab.

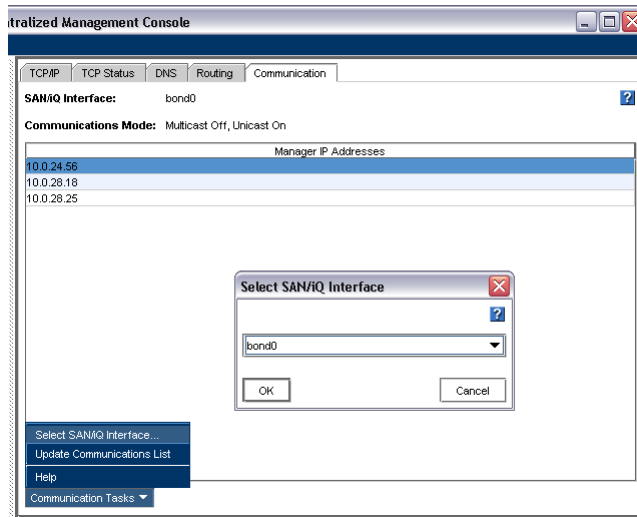


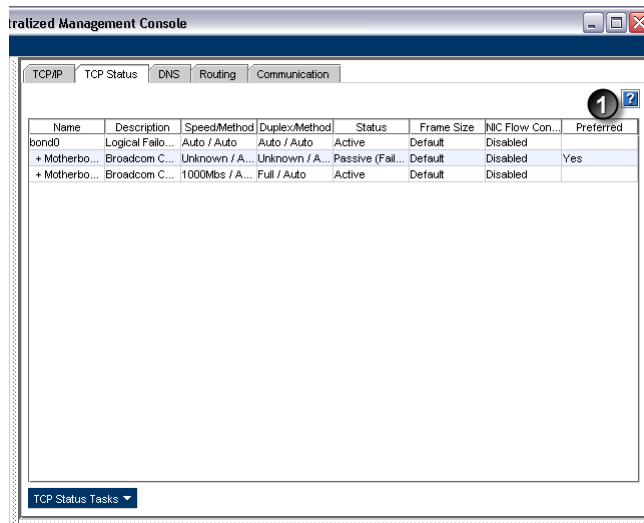
Figure 71 Verifying interface used for SAN/iQ communication

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

Viewing the Status of a NIC Bond

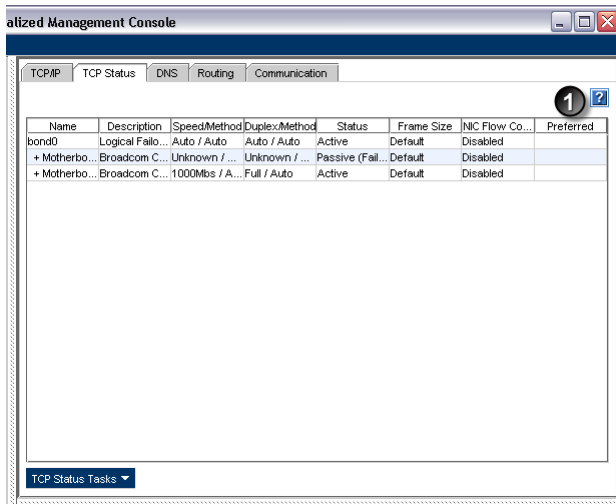
You can view the status of the interfaces on the TCP Status tab. Notice that in the Active-Passive bond, one of the NICs is the preferred NIC. In both the Link Aggregation Dynamic Mode bond and the Adaptive Load Balancing bond, neither physical interface is preferred.

Figure 72 shows the status of interfaces in an active-passive bond. Figure 73 shows the status of interfaces in a Link Aggregation Dynamic Mode bond.



1. Preferred interface

Figure 72 Viewing the status of an active-passive bond



1. Neither interface is preferred

Figure 73 Viewing the status of a link aggregation dynamic mode bond



NOTE:

If the bonded NIC experiences rapid, sequential Ethernet failures, the CMC may display the storage node as failed (flashing red) and access to data on that storage node fails. However, as soon as the Ethernet connection is reestablished, the storage node 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 node 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 that you want to delete.

4. Click on TCP/IP Tasks and select Delete Bond.

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

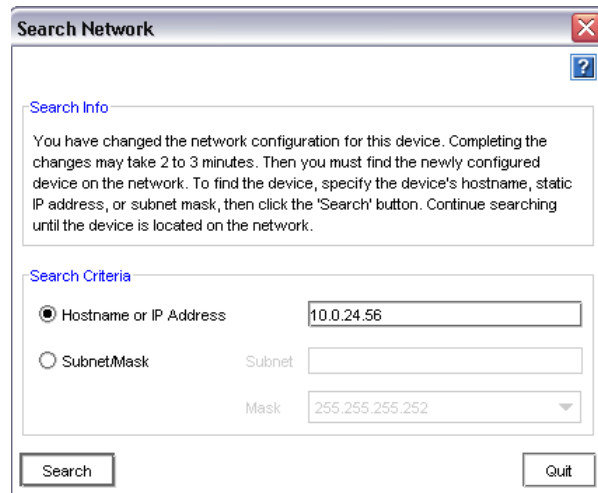


Figure 74 Searching for the unbonded storage node on the network

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



NOTE:

Because it can take a few minutes for the storage node 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](#)” on page 344.

Verify communication setting after deleting a bond

1. Select a storage node and open the tree below it.

2. Select the TCP/IP Network category and click the Communication tab.

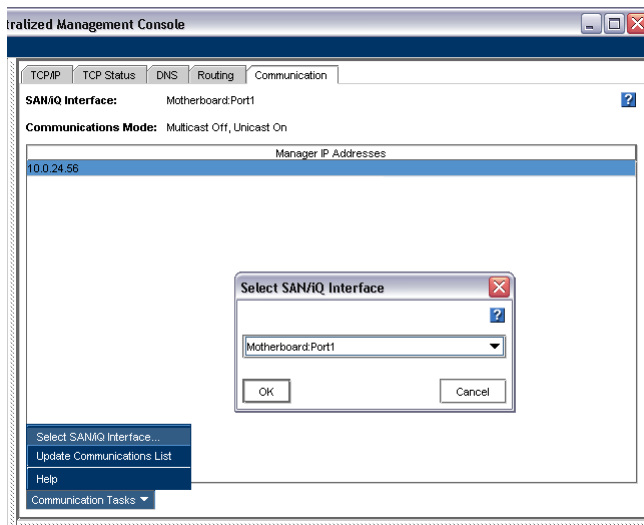


Figure 75 Verifying interface used for SAN/iQ communication

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

Disabling a network interface

You can disable the network interfaces on the storage node.

- 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 node, 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 node. If you disable all the interfaces, you must reconfigure at least one interface using the Configuration Interface to access the storage node. See [“Configuring a network connection”](#) on page 343.

Disabling a network interface

1. Log in to the storage node 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 node may be inaccessible if you continue.

7. Click OK.

If the storage node is in a management group

If the storage node 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 node but another interface is disconnected, you can reconnect to the second interface using the CMC. See “[Configuring the IP address manually](#)” on page 91.

If both interfaces to the storage node are disconnected, you must attach a terminal, or PC or laptop to the storage node with a null modem cable and configure at least one interface using the Configuration Interface. See “[Configuring a network connection](#)” on page 343.

TCP status

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.

The TCP status tab

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

Table 23 Status of and information about network interfaces

Column	Description
Name	<p>Name of the interface. Entries are</p> <ul style="list-style-type: none">• bond0—The bonded interface(s) (displays only if storage node is configured for bonding) <p>IBM x3650, NSM 160, Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300</p> <ul style="list-style-type: none">• Motherboard:Port1• Motherboard:Port2 <p>DL 380, DL 320s (NSM 2120)</p> <ul style="list-style-type: none">• G4-Motherboard:Port1• G4-Motherboard:Port2 <p>VSA</p> <ul style="list-style-type: none">• Eth0

Column	Description
	NSM 260 <ul style="list-style-type: none"> • Motherboard:Port0 • Motherboard:Port1
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 Table 16 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.

Managing settings on network interfaces

Configure or change the settings of the network interfaces in the storage nodes. See “[Network best practices](#)” on page 89 for more information.

Requirements

These settings must be configured before creating NIC bonds.

Changing speed and duplex settings

The settings for storage node and the switch must be the same. Available settings are listed in [Table 24](#).

Table 24 Setting storage node speed and duplex settings

Storage Node Setting	Speed/Duplex	Switch Setting	Speed/Duplex
Auto	Auto	Auto	Auto
1000	Full	1000	Full
100	Full	100	Full
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 node 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 node is in the Available Nodes pool and not in a management group.

To change the speed and duplex

1. In the navigation window, select the storage node 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 you want 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 displays. Then the changed setting displays 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”](#) on page 344.

Changing NIC frame size

Configure or change the settings of the network interfaces in the storage nodes. See [“Network best practices”](#) on page 89 for more information.

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 node is in the Available Nodes 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 node 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 nodes on the network, and set compatible frame sizes on all clients that access the storage nodes.

Consult with your network administrator for recommended storage node 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 node and also on each network switch between the storage node 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 node on the subnet must support jumbo frames.
- If you are using 802.1q virtual LANs, jumbo frames and non-jumbo 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 node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the TCP Status tab.
4. Select the interface you want to edit.
5. Click TCP Status Tasks and then 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 displays in the TCP status report.



NOTE:

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

Changing NIC flow control

You can 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 re-transmitted.



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 node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the TCP Status tab.
4. Select the interface you want to edit.
5. Click TCP Status Tasks and then select Edit.
6. Select On to enable the flow control on the NIC.
7. Click OK.
8. Repeat [Step 4](#) through [Step 7](#) for 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.

Using a DNS server

The storage node 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 node 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 node 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 storage node. These DNS servers are listed as IP addresses in the storage node configuration window in the TCP/IP Network category on the DNS tab. You can remove these DNS servers, but the storage node 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 node and you want the storage node to recognize host names, you must manually add a DNS server to the Network DNS tab.



NOTE:

If you initially set up the storage node 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 storage node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the DNS tab.

Adding the DNS domain name

Add the name of the DNS domain in which the storage node resides.

1. Click on DNS Tasks and then select Edit DNS Domain Name.
2. Type in 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 storage node.

1. Click on DNS Tasks and then select Edit DNS Server.
2. Click Add and type the IP address for the DNS server.
3. Click OK.
4. Repeat [Step 1](#) through [Step 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 node 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. Type the DNS suffix name. Use the domain name format.
4. Click OK.
5. Repeat [Step 1](#) through [Step 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 storage node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the DNS tab.
4. Select the server to edit.
5. Click DNS Tasks and select Edit DNS Servers.
6. Select the server again and click Edit.
7. Type 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 storage node.

1. In the navigation window, select a storage node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the DNS tab.
4. Click DNS Tasks and then select Edit DNS Domain Name.
5. Enter the change to the domain name.
6. Click OK.

Removing a DNS server

Remove a DNS server from the list.

1. In the navigation window, select a storage node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the DNS tab.
4. Select the server you want to remove from the DNS Servers list.
5. Click DNS Tasks and then select Edit DNS Servers.
6. Select the name again in the Edit DNS Servers window.
7. Click Remove.
8. 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 storage node and log in.
2. Open the tree and select the TCP/IP Network category.
3. Select the DNS tab.
4. Select the suffix you want to remove.
5. Click DNS Tasks and then select Edit DNS Suffixes.
6. Select the name again in the Edit DNS Suffixes window.
7. Click Remove.
8. 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 shut down of the storage node. To create a route that will survive a storage node reboot or shut down, you must enter a default gateway on the TCP/IP tab. See [“Configuring the IP address manually”](#) on page 91.

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

Adding routing information

1. In the navigation window, select a storage node 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. Type the IP address portion of the network address in the Net field.
8. Type 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 node 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 node 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 you want 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 node 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 you want to delete.
5. Click on Routing Tasks and select Edit Routing Information.
6. Select the routing information row you want to delete.
7. Click Delete.
8. Click OK on the confirmation message.

Configuring storage node communication

Use the Communication tab to configure the network interface used by the storage node to communicate with other storage nodes on the network and to update the list of managers that the storage node 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 nodes on the network. In order for clustering to work correctly, the SAN/iQ software communication interface must be designated on each storage node. The interface can be

- A single NIC that is not part of a bond
- A bonded interface consisting of 2 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 node 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 node 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.

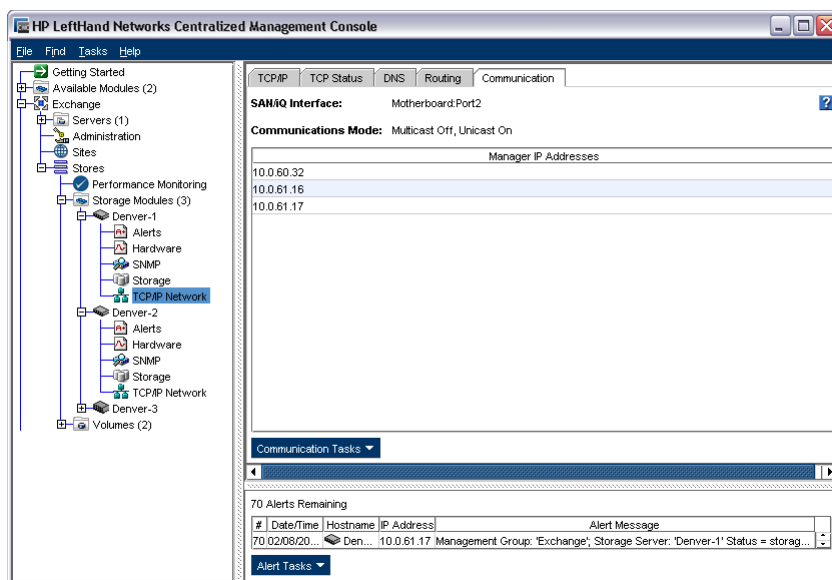


Figure 76 Selecting the SAN/iQ software network interface and updating the list of managers

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 node 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 node 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 node 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 node.

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

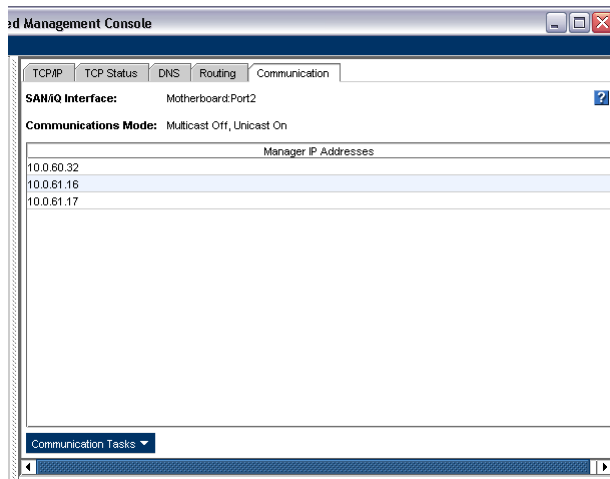


Figure 77 Viewing the list of manager IP addresses

4. Click Communication Tasks and select Update Communications List.

The list is updated with the current storage node 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 nodes 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 nodes inherit those management group settings.

- **Using network time protocol**

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

- **Setting the time zone**

Set the time zone for the storage node. 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 nodes 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 nodes 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 nodes 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 nodes 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 5 minute intervals. If you do not set the time zone for the management group, it uses GMT time zone.

 **NOTE:**

When using a Windows server as an external time source for an storage node, you must configure W32Time (the Windows Time service) to also use an external time source. The storage node 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. Type 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 node. **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 the NTP Server

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 type 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 the GMT time zone, whether or not you use NTP. Files display the time stamp according to this local time zone.

1. Click Time Tasks and then 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. Type a User Name and Description.
4. Type 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. Type 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 25](#). Users assigned to either of these groups assume the privileges associated with that group.

Table 25 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 node, 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. Type a Group Name and optional Description.

4. Select the permission level for each function for the group you are creating. See [Table 26](#).

Table 26 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 node. 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 Node 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 node.

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**—Users 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 node, such as read/write, and
- Access to different management capabilities for the storage node, 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 [Table 26](#) 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.

The Add Users window opens with a list of administrative users.

4. Select one or more 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. A confirmation window opens.

4. Click OK.

5. Click OK to finish.

7 Using SNMP

The storage node can be monitored using an SNMP Agent. You can also enable SNMP traps. The Management Information Base (MIB) is read-only and supports SNMP versions 1 and 2c. See [“Installing the LeftHand networks MIB”](#) on page 132 for a list of LeftHand Networks MIBs.

Using SNMP

SNMP and SNMP traps are disabled by default on all but the DL 320s (NSM 2120) and the DL 380. You must enable SNMP on the storage node and set a community string in order to access SNMP MIB objects.

Once you configure SNMP on a single storage node, you can copy the settings to other storage nodes using Copy Configuration. For more information, see [“Configuring multiple storage nodes”](#) on page 40.

SNMP on the DL 380 and DL 320s (NSM 2120)

SNMP is required for the HP server management system, Insight Manager. Insight Manager allows an administrator to manage and monitor a large number of servers from a single location. Because SNMP is required for Insight Manager, it is permanently enabled on the DL 380 and DL 320s (NSM 2120).

NOTE:

When you use the Insight Manger application with a DL 320s (NSM 2120), misleading attributes will display. Open Insight Manager and select File System Space Used. The display shows just one file system, which is not representative of the true space used.

Getting there

1. In the navigation window, select a storage node and log in.
2. Open the tree under the storage node and select the SNMP category.

Enabling SNMP agents

Most storage nodes allow enabling and disabling SNMP agents. On the DL 380 and the DL 320s (NSM 2120), SNMP is always enabled.

NOTE:

To query the LEFTHAND-NETWORKS-NSM-CLUSTER_MIB, you must have a manager running on the storage node servicing the SNMP request.

Adding an SNMP agent includes these tasks:

- Enabling the SNMP Agent
- Adding a community string.

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 management tool when attempting to access the system.

Community strings for DL 380 and DL 320s (NSM 2120)

Both the DL 380 and DL 320s (NSM 2120) have two reserved community strings that cannot be modified or removed.

- "Sanmon" is a read/write community
- "Public" is a read-only community.

If you use HP System Insight Manager with the DL 380 or DL 320s (NSM 2120), the user name and password for logging into the system are

- sanmon
- sanmon
- Adding access control for SNMP clients.

You can either enter a specific IP address and the IP Netmask as None to allow a specific host to access SNMP, or you can specify the Network Address with its netmask value so that all hosts matching that IP and netmask combination can access SNMP.

Additional agents and traps can be added and modified as on other storage nodes.



NOTE:

Use the CMC ping feature to verify IP addresses while configuring access control. See "[Pinging an IP address](#)" on page 91.

Enabling an SNMP agent

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.
3. On the SNMP General tab window, click SNMP General Tasks and select Edit SNMP Settings.
4. Select the Enabled radio button to activate the SNMP Agent fields.
5. Enter the Community String.
6. [Optional] Enter System Location information for the storage node.
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 network administrator information such as email address, or phone number for the person you would contact if you could not connect to SNMP clients.

Adding an SNMP client

- In the Access Control section, click Add to add an SNMP client that you can use to view SNMP. You can add SNMP clients by specifying either IP addresses or host names

By IP address

1. Select By Address and type the IP address.
2. Select an IP Netmask from the list. Select Single Host if adding only one SNMP client.
3. Click OK.
The IP address and netmask entry appear in the Access Control list.
4. Click OK in the Edit SNMP Settings window to finish.

By host name

1. Select By Name and type a host name.
That host name must exist in DNS and the NSM must be configured with DNS for the client to be recognized by the host name.
2. Click OK.
The host name appears in the Access Control list.
3. Click OK in the Edit SNMP Settings window to finish.

Editing access control entries

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.
3. On the SNMP General tab window, click SNMP General Tasks and select Edit SNMP 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 you are finished.

Deleting access control entries

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.
3. On the SNMP General tab window, click SNMP General Tasks and select Edit SNMP settings.
The Edit SNMP Settings window opens.
4. Select a client listed in the Access Control list and click Delete.
A confirmation message opens.
5. Click OK.

6. Click OK on the Edit SNMP Settings window when you are finished.

Using the SNMP MIB

The LeftHand Networks MIB provides read-only access to the storage node. The SNMP implementation in the storage node supports MIB-II compliant objects.

In addition, MIB files have been developed for storage node-specific information. These files, when loaded in the SNMP management tool, allow you to see storage node-specific information such as model number, serial number, hard disk capacity, network characteristics, RAID configuration, DNS server configuration details, and more.

Installing the LeftHand networks MIB

The LeftHand Networks MIB files are installed when you install the HP LeftHand Centralized Management Console if you selected a Complete installation. If you selected a Typical installation, you must load the LeftHand Networks MIB in the Network Management System as outlined below.

1. Load LEFTHAND-NETWORKS-GLOBAL-REG-MIB
2. Load LEFTHAND-NETWORKS-NSM-MIB
3. 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

This is the case when an SNMP management tool console and the HP LeftHand Centralized Management Console run on the same server.

NOTE:

Any variable that is labeled "Counter64" in the MIB requires version 2c or later of the protocol. Other standard version 2c compliant MIB files are also provided on the CD. Load these MIB files in the Network Management System if required.

Disabling the SNMP agent

Disable the SNMP Agent if you no longer want to use SNMP applications to monitor your network of storage nodes.

Disabling SNMP

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.

3. On the SNMP General tab window, click SNMP General Tasks and select Edit SNMP settings. The Edit SNMP Settings window opens.
4. Select Disable SNMP Agent.
5. Note that the Agent Status field shows disabled now. The SNMP client information remains listed, but cannot be used.

Adding SNMP traps

Prerequisite

- You must have first enabled an SNMP agent in order use SNMP traps. The DL 380 and the DL 320s (NSM 2120) always have an SNMP agent enabled.
- Add SNMP traps to have an SNMP tool send alerts when a monitoring threshold is reached.

Enable SNMP traps

You add a Trap Community String, used for client-side authentication.

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.
3. Select the SNMP Traps tab.
4. Click SNMP Trap Tasks and select Edit SNMP Traps.
5. Click Add to add trap recipients.
6. Enter the IP address or hostname for the SNMP client that is receiving the traps.
7. Click OK.
8. Repeat [Step 5](#) through [Step 7](#) for each host in the trap community.
9. Click OK on the Edit SNMP Traps window when you are finished adding hosts.

Editing trap recipients

1. Log in to the storage node and expand the tree.
2. Select the SNMP category from the tree.
3. Select the SNMP Traps tab.
The SNMP Traps Settings window opens.
4. Click SNMP Trap Tasks and select Edit SNMP Traps.
The Edit SNMP Traps window opens.
5. Select one of the Trap Recipients and click Edit.
6. Change the IP address or hostname and click OK.
7. Click OK when you are finished editing trap recipients.

Removing trap recipients

1. Log in to the storage node and expand the tree.

2. Select the SNMP category from the tree.
3. Select the SNMP Traps tab.
The SNMP Traps Settings window opens.
4. Click SNMP Trap Tasks and select Edit SNMP Traps.
The Edit SNMP Traps window opens.
5. Select one of the Trap Recipients and click Remove.
The host is removed from the list.
6. Click OK on the SNMP Traps tab when you are finished removing trap recipients.

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.

8 Reporting

Reporting capabilities of the HP LeftHand Storage Solution are divided into two categories:

- **Active Monitoring**—Use the Alerts category to configure how you receive alerts for selected variables. The Alerts category is where you set up email alerting and where you can review the log of alerts generated automatically by the operating system, including any generated while the CMC was closed. See [“Active monitoring overview”](#) on page 135.
- **Hardware Reporting**—Use the hardware category to select monitoring, perform hardware diagnostics, and generate a hardware report for the storage node. The Hardware category provides a report of system statistics, hardware, and configuration information. See [“Using the hardware information report”](#) on page 150.

Active monitoring overview

Alerts actively report on the condition of the hardware and storage network of a HP LeftHand Storage Solution. The Alerts category in the tree under every storage node includes multiple types of information and reporting capabilities. Review configuration information, save log files, set up email alerting, and review a log of alerts generated automatically by the operating system.

Use alerts to:

- View real-time statistical information about the storage node.
- View and save log files.
- Set up active monitoring of selected variables.
- Set up email notification.
- View alerts.

You can also set up SNMP traps to have an SNMP tool send alerts when a monitoring threshold is reached. For more information, see [“Adding SNMP traps”](#) on page 133.

Using alerts for active monitoring

Use active monitoring to track the health and operation of the storage node and management group. Active monitoring allows you to set up notification through emails, alerts in the CMC, and SNMP traps. You can choose which variables to monitor and choose the notification methods for alerts related to the monitored variables. Different storage nodes contain different sets of variables that can be monitored. For a detailed list of monitored variables, see [“List of monitored variables”](#) on page 138.



NOTE:

Critical variables, such as the Temperature Status (CPU and motherboard temperatures), have thresholds that trigger a shutdown of the storage node.

Getting there

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Alerts.

As you can see, some alerts are delivered to the console only, some include email delivery, and some are routed through the SNMP system as a trap.

Selecting alerts to monitor

When your software was first installed, all variables were selected to be reported on. You can change actively monitored variables as needed.

- Adding variables to monitor
See [“Adding variables to monitor”](#) on page 136.
- Removing variables from monitoring
See [“Removing a variable from active monitoring”](#) on page 138.
- Changing the way variables are monitored.
See [“Editing a monitored variable”](#) on page 137.

The section [“List of monitored variables”](#) on page 138 provides a list of all variables available for active monitoring.

Adding variables to monitor

The variables that the storage node is currently monitoring are listed in the box. All variables in the list are configured and set for CMC alerts.

1. Click the Alert Setup tab to bring it to the front.
2. Click Alert Setup Tasks and select Add Monitored Variables.
3. Select the variable that you want to begin to monitor and click Next.
4. Specify the frequency for monitoring the variable and click Next.

5. For each threshold listed, select the type of alert you want to receive.

Table 27 Types of alerts for active monitoring

Type of alert	Where alerts are sent	For more information
CMC alerts	To the alert window of the CMC and the Alerts tab in Reporting.	See “Using the alert window” on page 34.
SNMP traps	To the SNMP trap community managers. You must have configured the storage node to use SNMP.	See “Adding SNMP traps” on page 133.
Email	To specified email addresses. Type the email addresses to receive the notification, separated by commas. Then configure Email Notification on the Email tab.	See “Adding variables to monitor” on page 136.

 **NOTE:**

To save time while setting up active monitoring, select all variables, then click Set Notifications. This setting applies the same email address and other alert settings to all storage nodes. Then, if you need to customize alert actions for a particular variable, you can edit that variable.

6. Click Finish when you have configured all the threshold items in the list.

Editing a monitored variable

For the selected variable, you can change the frequency of monitoring and the notification routing of the alert.

1. Select the Alert Setup tab.
2. Select the variable you want to edit.
3. Click Alert Setup Tasks and select Edit Monitored Variable.

The Configure Variable wizard opens to Step 1.

 **NOTE:**

For some variables, only the notification method can be changed. For example, the frequency for the Storage Server Latency variable is set to 1 minute and cannot be changed.

4. If allowed, change the frequency for the variable and click Next.
The Configure Variable wizard opens to Step 2.
5. (Optional) Change the alert notification method.
6. Click Finish.

 **NOTE:**

If you are requesting email notification, be sure to set up the SMTP settings on the Email Server Setup tab.

Removing a variable from active monitoring

Use Remove to remove variables to stop active monitoring that become pointless or impractical. You can return a variable to active monitoring at any time. Permanent variables, such as Cache Status, cannot be removed. See “[List of monitored variables](#)” on page 138.

1. Show the Alert Setup tab window.
2. Select the variable you want to remove.
3. Click Alert Setup Tasks and select Remove Monitored Variable.
4. Click Remove.
A confirmation message opens.
5. Click OK in the confirm window.
The variable is removed.

NOTE:

Variables are not deleted when they are removed from active monitoring. You can add them back to active monitoring at any time.

List of monitored variables

This section contains tables that show the variables that are monitored during active, not passive, monitoring. For each variable, the table lists the following information:

- The units of measurement.
- Whether the variable is permanent. (Permanent variables cannot be removed from active reporting.)
- Whether you can change the frequency with which the measurements are taken.
- The default frequency of measurements.
- The default action that occurs if the measured value of the variable reaches a threshold.

Table 28 List of monitored variables

Variable name	Units/ status	Perm. variable	Specify freq.	Default freq.	Default action/threshold
BBU Capacity Test (24Hr) [NSM 160] ¹	Status	Yes	Yes	Saturday, 23:00 Monthly	CMC alert if failed
BBU Capacity Test Overdue [NSM 160]	Status	Yes	Yes	1 hour	CMC alert if changes
Boot Device Status [NSM 160, NSM 260, NSM 4150]	Status	No	Yes	1 minute	CMC alert if changes
CPU Utilization	Percent	No	Yes	1 minute	None

Variable name	Units/ status	Perm. variable	Specify freq.	Default freq.	Default action/threshold
Cache Status	Status	Yes	Yes	1 minute	CMC alert if changes
Cluster Utilization	Percent	Yes	Yes	15 minutes	CMC alert if the value exceeds 90 CMC alert if the value exceeds 95
Cluster Virtual IP Status	Normal, Faulty	No	Yes	1 hour	CMC alert if not Normal
Drive Health	Status	Yes	Yes	1 minute	CMC alert if change or critical
Drive Status NSM 160: 1 through 4 NSM 260: 1 through 12 DL 380: 0 through 5 DL 320s: 1 through 12 Dell 2950: 0 through 5 NSM 2060: 0 through 5 NSM 4150: 0 through 14 HP LeftHand P4500: 1 through 12 HP LeftHand P4300: 1 through 8	Status	No	Yes	1 minute	CMC alert if changes
Fan Status	Status	Yes	Yes	1 minute	CMC alert if not Normal
LogPart Utilization	Percent	Yes	Yes	2 minutes	CMC alert if the value exceeds 95 CMC alert if the value exceeds 80
Management Group Maintenance Mode	True, False	Yes	Yes	15 minutes	CMC alert if True
Memory Utilization	Percent	No	Yes	1 minute	CMC alert at > 90%

Variable name	Units/ status	Perm. variable	Specify freq.	Default freq.	Default action/threshold
Network Interface Status	Status	No	Yes	1 minute	CMC alert if NIC status changes
Power Supply Status	Status	No	Yes	1 minute	CMC alert if status changes
RAID Status	Status	Yes	Yes	15 seconds	CMC alert if changes
Remote Copy Complete	True, False	No	Yes	15 minutes	CMC alert when True
Remote Copy Failover	True, False	No	Yes	15 minutes	CMC alert when True
Remote Copy Status	Status	No	Yes	15 minutes	CMC alert if not Normal
Remote Management Group Status	Up, Down	No	Yes	1 minute	CMC alert if changes
SAN/iQ Memory Requirement	Status	Yes	No	1 minute	CMC alert if fails
Snapshot Schedule Status	Status	No	Yes	1 minute	CMC alert if snapshot status is not Normal
Storage Server Latency	Milliseconds	Yes	No	1 minute	CMC alert if > 60 seconds
Storage Server Status	Up, Down	No	Yes	1 minute	CMC alert if not Up
Temperature Status	Status	Yes	No	1 minute	CMC alert if warning level reached CMC alert and shutdown: if critical level reached See the alert or the Hardware Information tab in the CMC for additional temperature information.
Voltage Status [NSM 160, NSM 260]	Status	Yes	Yes	1 minute	CMC alert if not Normal

Variable name	Units/ status	Perm. variable	Specify freq.	Default freq.	Default action/threshold
Volume Re-stripe Complete	True, False	No	Yes	1 minute	CMC alert when True
Volume Status	Status	No	Yes	1 minute	CMC alert if volume status changes
Volume Threshold Change	Status	Yes	Yes	1 minute	CMC alert when True
Volume Thresholds	Status	No	Yes	15 minutes	CMC alert if threshold exceeded for any volume or snapshot in the management group

¹BBU capacity test runs monthly to monitor the battery life (charge). If the battery life remaining is less than 72 hours, the cache is shut off to protect data. When the cache is shut off, there will be a decrease in performance.

Setting CMC notification of alerts

By default, you will see all alerts displayed on the console in the bottom Alerts window. This method of notification may be turned off and on.

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Alerts.
3. Select the Alert Setup tab.
4. Select any alert in the list.
5. Click alert Setup Tasks and select Set Threshold Actions.
6. Click the checkbox for HP LeftHand Centralized Management Console alert.
7. Click OK.

Setting SNMP notification of alerts

Configure the following items to have alerts delivered via SNMP.

- Enable the SNMP system.
See [“Enabling SNMP agents”](#) on page 129.
- Select the check box for SNMP Trap.
See either [“Setting notification for one variable”](#) on page 142 or [“Setting notification for several variables”](#) on page 143

Setting email notifications of alerts

If you request email notification of alerts, you must configure these settings.

- SMTP settings
See [“Setting SMTP settings”](#) on page 142
- Application of SMTP settings to an entire management group
See [“Applying SMTP settings to a management group”](#) on page 142
- Email notification preferences
See either [“Setting notification for one variable”](#) on page 142 or [“Setting notification for several variables”](#) on page 143

Setting SMTP settings

Use the Email Server Setup tab to configure the SMTP settings for email communication. For more information on configuring active monitoring, see [“Using alerts for active monitoring”](#) on page 135.

1. In the Alerts category, select the Email Server Setup tab.
The Email Server Setup window opens.
2. Click Email Server Setup Tasks and select Edit SMTP Settings.
3. Enter the IP address or host name of the email server.
4. Enter the email port.
The standard port is 25.
5. (Optional) If your email server is selective about valid sender addresses on incoming emails, enter a sender address, for example, “username@company.com.”
If you do not enter a sender address, the From field of email notifications will display “root@hostname,” where hostname is the name of the storage node.
6. Select the box if you want to apply the settings to all the storage nodes in the management group.
If you DO apply the settings to other storage nodes in the management group, and if you have a sender address entered, all the other storage nodes will use that same sender address.
7. (Optional) Test the email connectivity now if you wish.
8. Click OK.

NOTE:

Notification of undeliverable email messages are sent to the sender address.

If you are requesting email notification, be sure to set up the email notification in Alert Setup.

Applying SMTP settings to a management group

Perform the steps in [“Setting SMTP settings”](#) on page 142, and check the check box to Apply these SMTP settings to all storage nodes in the management group.

Setting notification for one variable

Specify email notification in two ways:

- By editing an individual monitored variable.
See [“Editing a monitored variable”](#) on page 137.

- By setting the Set Threshold Actions window and affecting several monitored variables.

Setting notification for several variables

Setting threshold actions determines the routing preferences for notification of alerts. With this procedure, you can set the same email address for several alerts.

1. Select the Alert Setup tab.
2. Select the variables you want to change.
3. Click Alert Setup Tasks and select Set Threshold Actions.
4. Click the check boxes to specify where you want the alert to be communicated.
 - If you check SNMP Trap, you must enable SNMP in that category of the storage node. See “[Enabling SNMP agents](#)” on page 129.
 - If you check Email, you must provide an email address to send to, and you must set up SMTP addresses as well.
5. Click OK.

Viewing and saving alerts

Any time that an actively monitored variable causes an alert, the alert is logged by the storage node. If the CMC is open, alerts display in the Alert window on the CMC main window.

If the CMC is not open, these alerts are still logged, and you can view them the next time you open the CMC. Click a storage node > Alerts > Alert Log File tab.

NOTE:

Alerts category > Alerts Log File under a storage node displays the most recent alerts, up until the alert list reaches 1 MB in size. To view alerts older than those displayed on the Alerts tab, save the Alerts log on the Alert Log Files tab.

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Alerts.
3. Select the Alert Log File tab.
4. Click Refresh to make sure you view the most current data.

Saving the alert log of all variables

1. Perform the tasks in “[Viewing and saving alerts](#)” on page 143.
That is, select a storage node > Alerts > Alert Log Files tab.
2. To save the list of alerts, click Alert Log File Tasks and select Save to File.
3. Select a location for the file.
This file will never be more than one megabyte.

Saving the alert history of a specific variable

To save the history of a specific variable on a specific storage node, save a copy of the log file for that variable. This copy is a text file with file name the same as the variable.

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Alerts.
3. Select the Alert Setup tab.
4. Highlight the variable for which you want to save the log file.
This selects it. Use CTRL and click to select several variables. A separate file is created for each.
5. In the Alert Setup Task pull down menu, select Save Log Files.
A Save window opens.
6. Choose a location for the file.
7. Click Save.
The file is saved to the location you specified. Use a file manager window and text editor to check it.

Using hardware information reports

The Hardware category, found in the tree under every storage node, includes multiple types of information and reporting capabilities. Review a Hardware report of system statistics, hardware, and configuration information.

Use the Hardware category to:

- Run hardware diagnostics. See [“Running diagnostic reports”](#) on page 144.
- View storage node hardware information in real-time. See [“Using the hardware information report”](#) on page 150.
- View and save a storage node’s log files. See [“Saving log files”](#) on page 167.
- View and save log files to a remote computer. See [“Using hardware information log files”](#) on page 167.

Running diagnostic reports

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



NOTE:

Running diagnostics can help you to monitor the health of the storage node or to troubleshoot hardware problems.

Getting there

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Hardware.

3. Select from the list the diagnostic tests that you want to run.

The default setting is to run all tests. If all boxes are not checked, click Diagnostic Tasks and select Check All. Clear any tests that you do not want to run. To clear all selections, click Clear All.

 **NOTE:**

Running all of the diagnostic tests will take several minutes. To shorten the time required to run tests, clear the checkboxes for any tests that you do not need.

4. Click Diagnostic Tasks and select Run Tests.
A progress message displays. When the tests complete, the results of each test display in the Result column.
5. (Optional) When the tests complete, if you want to view a report of test results, click Save to File. Then select a location for the diagnostic report file and click Save.
The diagnostic report is saved as a .txt file in the designated location.

Viewing the diagnostic report

The results of diagnostic tests are written to a report file. For each diagnostic test, the report lists whether the test was run and whether the test passed, failed, or issued a warning.

 **NOTE:**

If any of the diagnostics show a result of "Failed," call Customer Support.

To view the report file:

1. After the diagnostic tests complete, save the report to a file.
2. Browse to the location where you saved the diagnostics report (.txt) file.
3. Open the report file.

List of diagnostic tests

This section shows the diagnostic tests that are available for the storage node. For each test, the table lists the following information:

- A description of the test
- Pass / fail criteria

See the specific table for your platform.

- For NSM 160 and NSM 260, see [Table 29](#)
- For DL 380, DL 320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300, see [Table 30](#)
- For IBM x3650, see [Table 31](#)

- For Dell 2950, NSM 2060, and NSM 4150, see [Table 33](#)

Table 29 List of hardware diagnostic tests and pass / fail criteria for NSM 160 and NSM 260

Diagnostic Test	Description	Pass criteria	Fail criteria	NSM 160	NSM 260
Fan Test	Checks the status of all fans.	Fan is normal	Fan is faulty or missing	X	X
Power Test	Checks the status of all power supplies.	Supply is normal	Supply is faulty or missing	X	X
Temperature Test	Checks the status of all temperature sensors.	Temperature is within normal operating range	Temperature is outside normal operating range	X	X
Voltage Test	Checks the status of all voltage sensors.	Voltage is within normal operating range	Voltage is outside normal operating range	X	X
Cache Status	Checks the status of the disk controller caches.	Cache is normal	Cache is corrupt	X	X
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	X	X
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing	X	X
Disk Temperature Test	Checks the temperature of all disk drives.	The temperature is within normal operating range	The temperature is outside normal operating range	X	X
Disk SMART Health Test	S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks. A program inside the disk constantly tracks a range of the vital characteristics, including driver, disk heads, surface state, and electronics. This information may help predict hard drive failures.	All drives pass health test	Warning or Failed if one or more drives fails health test	X	X
Generate SMART logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	X	X

Diagnostic Test	Description	Pass criteria	Fail criteria	NSM 160	NSM 260
Generate 3ware Diagnostics Report (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	X	-
BBU Capacity Test	Tests the ability of the BBU to hold a charge. BBUs weaken over time. A failure indicates it is time to replace the BBU.	The BBU can hold an acceptable charge	The BBU failed to hold an acceptable charge	X	-

Table 30 List of hardware diagnostic tests and pass/fail criteria for DL 380, DL 320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300

Diagnostic Test	Description	Pass criteria	Fail criteria	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
Fan Test	Checks the status of all fans.	Fan is normal	Fan is faulty or missing	X	X	X	X
Power Test	Checks the status of all power supplies.	Supply is normal	Supply is faulty or missing	X	X	X	X
Temperature Test	Checks the status of all temperature sensors.	Temperature is within normal operating range	Temperature is outside normal operating range	X	X	X	X
Cache Status	Checks the status of the disk controller caches.	Cache is normal	Cache is corrupt	X	X	X	X
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	X	X	X	X
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing	X	X	X	X
Disk Temperature Test	Checks the temperature of all disk drives.	The temperature is within normal operating range	The temperature is outside normal operating range	X	-	-	-

Diagnostic Test	Description	Pass criteria	Fail criteria	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
Disk SMART Health Test	S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks. A program inside the disk constantly tracks a range of the vital characteristics, including driver, disk heads, surface state, and electronics. This information may help predict hard drive failures.	All drives pass health test	Warning or Failed if one or more drives fails health test	X	X	X	X
Generate SMART logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	X	-	X	X
Generate HP Diagnostic Report (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	-	X	X	X

Table 31 List of hardware diagnostic tests and pass/fail criteria for IBM x3650

Diagnostic Test	Description	Pass criteria	Fail criteria	IBM x3650
Fan Test	Checks the status of all fans.	Fan is normal	Fan is faulty or missing	X
Power Test	Checks the status of all power supplies.	Supply is normal	Supply is faulty or missing	X
Temperature Test	Checks the status of all temperature sensors.	Temperature is within normal operating range	Temperature is outside normal operating range	X
Cache Status	Checks the status of the disk controller caches.	Cache is normal	Cache is corrupt	X
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	X

Diagnostic Test	Description	Pass criteria	Fail criteria	IBM x3650
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing	X
Disk Temperature Test	Checks the temperature of all disk drives.	The temperature is within normal operating range	The temperature is outside normal operating range	X
Disk SMART Health Test	S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks. A program inside the disk constantly tracks a range of the vital characteristics, including driver, disk heads, surface state, and electronics. This information may help predict hard drive failures.	All drives pass health test	Warning or Failed if one or more drives fails health test	X
Generate SMART logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	X
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	X

Table 32 List of hardware diagnostic tests and pass/fail criteria for VSA

Diagnostic Test	Description	Pass criteria	Fail criteria	VSA
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing	X

Table 33 List of hardware diagnostic tests and pass/fail criteria for Dell 2950, NSM 2060, and NSM 4150

Diagnostic Test	Description	Pass criteria	Fail criteria	Dell 2950	NSM 2060	NSM 4150
Fan Test	Checks the status of all fans.	Fan is normal	Fan is faulty or missing	X	X	X
Power Test	Checks the status of all power supplies.	Supply is normal	Supply is faulty or missing	X	X	X
Temperature Test	Checks the status of all temperature sensors.	Temperature is within normal operating range	Temperature is outside normal operating range	X	X	X
Cache Status	Checks the status of the disk controller caches.	Cache is normal	Cache is corrupt	X	X	X

Diagnostic Test	Description	Pass criteria	Fail criteria	Dell 2950	NSM 2060	NSM 4150
Cache BBU Status	Checks the status of the Battery Backup Unit (BBU).	The BBU is normal and not charging or testing	The BBU is charging, testing or faulty	X	X	X
Disk Status Test	Checks for the presence of all disk drives.	All disk drives are present	One or more drives are missing	X	X	X
Disk SMART Health Test	S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology) is implemented in all modern disks. A program inside the disk constantly tracks a range of the vital characteristics, including driver, disk heads, surface state, and electronics. This information may help predict hard drive failures.	All drives pass health test	Warning or Failed if one or more drives fails health test	X	X	X
Generate SMART logs (for analysis contact Customer Support)	Generates a drive health report.	The report was successfully generated	The report was not generated	X	X	X
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	X	X	X

Using the hardware information report

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

Generating a hardware information report

To generate a Hardware Information report

1. Select the Hardware Information tab.

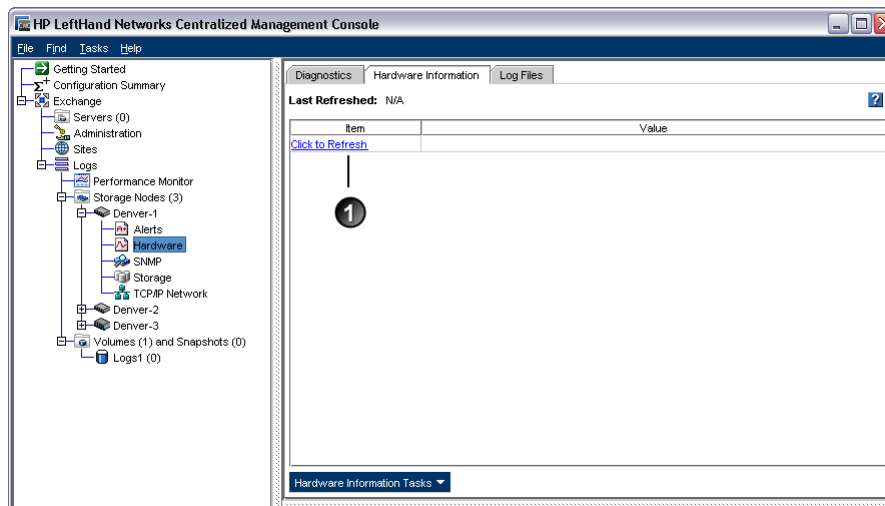


Figure 78 Opening the hardware information window

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

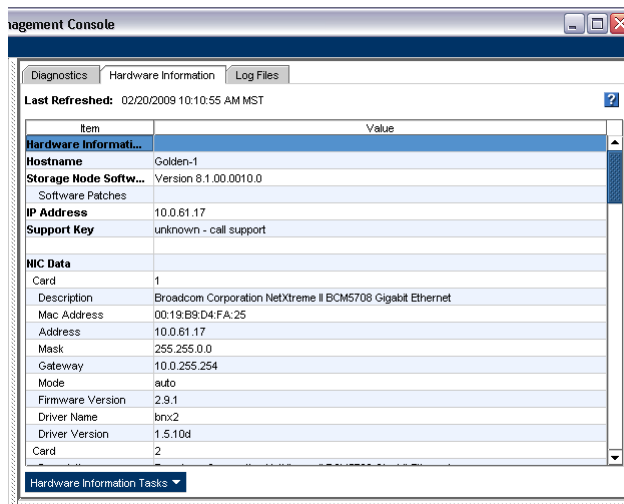


Figure 79 Viewing the hardware information for a storage node

Saving a hardware information report

1. Click Hardware Information Tasks and select Save to File to download a text file of the reported statistics.
The Save dialog opens.
2. Choose the location and name for the report.
3. Click Save.
The report is saved with an .html extension.

Hardware information report details

This section contains more information about the Hardware Information report for the following platforms:

- [Table 34](#) on page 152
- [Table 35](#) on page 157
- [Table 36](#) on page 162
- [Table 37](#) on page 163

Table 34 Selected details of the Hardware report for the NSM 160, NSM 260, and VSA

This term	Means This	NSM 160	NSM 260	VSA
Last refreshed	Date and time report created.	X	X	--
Hardware information	Date and time report created.	X	--	--
Name or hostname	Hostname of the storage node.	X	X	X
IP address	IP address of the storage node.	X	X	X
Storage node software	Full version number for storage node software. Also lists any patches that have been applied to the storage node.	X	X	X
Support Key	Support Key is used by a Technical Support representative to log in to the storage node. (Not available for the Demo version.)	X	X	X

This term	Means This	NSM 160	NSM 260	VSA
NIC data	Information about NICs in the storage node, including the card number, manufacturer description, MAC address, IP address, mask, gateway, and mode. Mode shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, and disabled means the interface is disabled.	X	X	X
DNS data	Information about DNS, if a DNS server is being used, providing the IP address of the DNS servers.	X	X	X
Memory	Information about RAM memory in the storage node, including, as necessary, the values for total, free, shared, cached memory, and number of buffers.	X	X	X
CPU	This section of the report contains many details about the CPU, including model name, clock speed, and cache size.	X	X	--

This term	Means This	NSM 160	NSM 260	VSA
Stat	<p>Information about the CPU.</p> <ul style="list-style-type: none"> • CPU seconds shows the number of CPU seconds spent on user tasks, kernel tasks, and in idle state. • Machine up-time is the total time the storage node has been running from initial boot up. 	X	X	X
Backplane information	<p>This part of the report delivers selected information about the backplane, firmware version, serial number, and LED information.</p>	--	X	--
Motherboard information	<p>This part of the report delivers selected information about the motherboard, including, but not limited to, IPMI, and firmware.</p>	X	X	--
Drive status	<p>For each drive, reports the status, health, and temperature. [VSA only] Temperature is not reported for the VSA. Health = normal if a drive is present and powered on.</p>	--	X	X
Drive Info	<p>For each drive, reports the model, serial number, and capacity.</p>	--	X	X
RAID	<p>Information about RAID.</p>	X	--	X

This term	Means This	NSM 160	NSM 260	VSA
Rebuild Rate	RAID Rebuild Rate is a percentage of RAID card throughput.	--	X	--
	RAID Rebuild Rate is a priority measured against other OS tasks.	X	--	--
Unused Devices	<p>Any device which is not participating in RAID. This includes:</p> <ul style="list-style-type: none"> • Drives that are missing • Unconfigured drives • Drives that were powered down • Failed drives (rejected by array w/IO errors) • Drives that are rebuilding • Hot-spare drives 	X	X	X
Statistics	Information about the RAID for the storage node.	X	X	X
Unit Number	<p>Identifies devices that make up the RAID configuration, including:</p> <ul style="list-style-type: none"> • Type of storage (BOOT, LOG, SANIQ, DATA) • RAID level (0, 1, 5, Virtual) • Status (Normal, Rebuilding, Degraded, Off) • Capacity • Rebuild statistics (% complete, time remaining) 	X	X	X
RAID O/S partitions	Information about O/S RAID.	--	X	X

This term	Means This	NSM 160	NSM 260	VSA
Minimum Rebuild Speed	Minimum amount of data in MB/seconds that will be transferred during an O/S RAID rebuild. The higher this number, the less bandwidth available for users because the system will not transfer at a rate lower than what is set.	X	X	X
Maximum Rebuild Speed	The maximum amount of data in MB/second that will be transferred during an O/S RAID rebuild.	X	X	X
Statistics	Information about the O/S RAID for the storage node.	X	X	X
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) 	X	X	X
Boot device statistics	Disk number, flash status, capacity, driver version, media used for device, and model number.	--	X	--
Power supply	Information about the type of power supplies in the storage node.	X	X	--

This term	Means This	NSM 160	NSM 260	VSA
Power supplies	Status information about those power supplies.	X	X	--
Controller cache items	Information about RAM, including but not limited to the model, serial number, status, battery status, versioning, cache size, memory size, and voltage.	X	X	X
Sensor data	For the hardware listed, shows information about fan, voltage and temperature sensors on the mother board, including minimums and maximum.	X	X	--

Table 35 Selected details of the Hardware Report for DL 380, DL 320s (NSM 2120), HP LeftHand P4500, and HP LeftHand P4300

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
Last refreshed	Date and time report created.	X	X	X	X
Hostname	Hostname of the storage node.	X	X	X	X
Storage node software	Full version number for storage node software. Also lists any patches that have been applied to the storage node.	X	X	X	X
IP address	IP address of the storage node.	X	X	X	X
Support key	Support Key is used by a Technical Support representative to log in to the storage node.	X	X	X	X

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
NIC data	Information about NICs in the storage node: card number, manufacturer description, MAC address, IP address, mask, gateway, and mode. Mode shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, and disabled means the interface is disabled.	X	X	X	X
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.	X	X	X	X
Memory	Information about RAM in the storage node, including values for total memory and free memory in GB.	X	X	X	X
CPU	This section of the report contains details about the CPU, including model name or manufacturer of the CPU, clock speed of the CPU, and cache size.	X	X	X	X

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
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 node has been running from initial boot up.	X	X	X	X
Backplane information	This part of the report delivers selected information about the backplane LEDs: LED support and idLED.	X	X	X	X
Drive info	For each drive, reports the model, serial number, and capacity.	X	X	X	X
Drive status	For each drive, reports the status, health, and temperature.	X	X	X	X
RAID	Information about RAID.	X	X	X	X
Rebuild rate	RAID Rebuild Rate is a priority measured against other OS tasks.	X	X	X	X

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
Unused devices	<p>Any device which is not participating in RAID. This includes:</p> <ul style="list-style-type: none"> • Drives that are missing • Unconfigured drives • Drives that were powered down • Failed drives (rejected by array w/IO errors) • Drives that are rebuilding • Hot-spare drives 	X	X	X	X
Statistics	Information about the RAID for the storage node.	X	X	X	X
Unit number	<p>Identifies devices that make up the RAID configuration, including:</p> <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) 	X	X	X	X

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
RAID O/S partitions	Information about O/S RAID.	X	X	X	X
Statistics	Information about the O/S RAID for the storage node.	X	X	X	X
Unit number	<p>Identifies devices that make up the O/S RAID configuration, including:</p> <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) 	X	X	X	X
Controller/cache items	Information about the RAID controller card and Battery Backup Unit (BBU) including: model number, serial number, cache status, battery status, hardware version, and firmware version.	X	X	X	X
Power supply	Shows the type or number of power supplies.	X	X	X	X

This term	Means this	DL 380	DL 320s	HP LeftHand P4500	HP LeftHand P4300
Power supplies	Status information about the power supplies.	X	X	X	X
Sensors	Shows for the hardware listed, the status, real measured value, minimum and maximum values.	X	X	X	X

Table 36 Selected details of the hardware report for IBM x3650

This term	Means this	IBM x3650
Last refreshed	Date and time report created.	X
Hostname	Hostname of the storage node.	X
IP number	IP address of the storage node.	X
Storage node software	Full version number for storage node software. Also lists any patches that have been applied to the storage node.	X
Support key	Support Key is used by a Technical Support representative to log in to the storage node.	X
NIC data	Information about NICs in the storage node including the card number, manufacturer description, MAC address, mask, gateway, and mode. Mode shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, and disabled means the interface is disabled.	X
DNS data	Information about DNS, if a DNS server is being used, providing the IP address of the DNS servers.	X
Memory	Information about RAM memory in the storage node: the total amount of memory in GB, and the total amount of free memory in GB.	X

This term	Means this	IBM x3650
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 node has been running from initial boot up.	X
Drive info	For each drive, reports the model, serial number, and capacity of the drive.	X
Drive status	For each drive, reports the status, health, and temperature.	X
RAID	Information about RAID.	X
Rebuild rate	RAID Rebuild Rate is a percentage of RAID card throughput.	X
Unused devices	Any device which is not participating in RAID. This includes: <ul style="list-style-type: none"> • Drives that are missing • Unconfigured drives • Drives that were powered down • Failed drives (rejected by array w/IO errors) • Drives that are rebuilding • Hot-spare drives 	X
Statistics	Information about the RAID for the storage node.	X
Unit number	Identifies devices that make up the 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) 	X
Sensors	Does not apply to this platform.	--

Table 37 Selected details of the hardware Report for Dell 2950, NSM 2060, and NSM 4150

This term	Means this	Dell 2950	NSM 2060	NSM 4150
Last refreshed	Date and time report created.	X	X	X

This term	Means this	Dell 2950	NSM 2060	NSM 4150
Hostname	Hostname of the storage node.	X	X	X
Storage node software	Full version number for storage node software. Also lists any patches that have been applied to the storage node.	X	X	x
IP address	IP address of the storage node.	X	X	X
Support key	Support Key is used by a Technical Support representative to log in to the storage node.	X	X	X
NIC data	Information about NICs in the storage node: card number, manufacturer description, MAC address, IP address, mask, gateway, and mode. Mode shows manual/auto/disabled. Manual equals a static IP, auto equals DHCP, and disabled means the interface is disabled.	X	X	X
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.	X	X	X
Memory	Information about RAM in the storage node, including values for total memory and free memory in GB.	X	X	X

This term	Means this	Dell 2950	NSM 2060	NSM 4150
CPU	This section of the report contains details about the CPU, including model name or manufacturer of the CPU, clock speed of the CPU, and cache size.	X	X	X
Stat	Information about the CPU. CPU seconds shows the number of CPU seconds spent on user tasks, kernel tasks, and in idle state. Machine up-time is the total time the storage node has been running from initial boot up.	X	X	X
Motherboard information	Serial number of the chassis.	X	X	X
Drive info	For each drive, reports the model, serial number, capacity, and firmware version.	X	X	X
Drive status	For each drive, reports the status and health.	X	X	X
RAID	Information about RAID.	X	X	X
Rebuild rate	RAID Rebuild Rate is a priority measured against other OS tasks.	X	X	X

This term	Means this	Dell 2950	NSM 2060	NSM 4150
Unused devices	<p>Any device which is not participating in RAID. This includes:</p> <ul style="list-style-type: none"> • Drives that are missing • Unconfigured drives • Drives that were powered down • Failed drives (rejected by array w/IO errors) • Drives that are rebuilding • Hot-spare drives 	X	X	X
Statistics	Information about the RAID for the storage node.	X	X	X
Unit number	<p>Identifies devices that make up the RAID configuration, including:</p> <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) 	X	X	X
RAID O/S partitions	Information about O/S RAID.	X	X	X
Minimum rebuild speed	Minimum speed at which the system can rebuild data.	X	X	X
Maximum rebuild speed	Maximum speed at which the system can rebuild data.	X	X	X

This term	Means this	Dell 2950	NSM 2060	NSM 4150
Statistics	Information about the O/S RAID for the storage node.	X	X	X
Boot-device statistics	Status information about the boot device: status, capacity in MB, driver version, media used for device, and model.	-	-	X
Controller/cache items	Information about the RAID controller card and Battery Backup Unit (BBU) including: model number, serial number, cache status, battery status, hardware version, and firmware version.	X	X	X
Power supply	Shows the type or number of power supplies.	X	X	X
Power supplies	Status information about the power supplies.	X	X	X
Sensors	Shows for the hardware listed, the status, real measured value, minimum and maximum values.	X	X	X

Using hardware information log files

The log files that contain hardware information are always saved on the individual storage node. You may want to save those log files to another computer. This way, if the storage node goes off line, the log files will be available.

This section explains how to save a hardware information log file to a txt file on the local storage node or a remote computer. See these sections:

- [“Saving log files”](#) on page 167
- [“Using remote log files”](#) on page 168

Saving 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 node (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).

 **NOTE:**

Save the log files that are stored locally on the storage node. For more information about remote log files, see [“Using remote log files”](#) on page 168.

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Hardware.
3. Select the Log Files tab.
4. To make sure you have the latest data, click Log File Tasks and select Refresh Log File List.
5. Scroll down the list of Choose Logs to Save and select the file or files you want to save.
To select multiple files, use the Ctrl key.
6. Click Log Files Tasks and select Save Log Files.
The Save window opens.
7. Select a location for the file or files.
8. Click Save.
The file or files are saved to the designated location.

Using remote log files

Use remote log files to automatically write log files to a computer other than the storage node. For example, you can direct the log files for one or more storage nodes 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.

Adding a remote log

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Hardware.
3. Select the Log Files tab.
4. Click Log File Tasks and select Add Remote Log Destination.
5. In the Log Type drop-down list, select the log you want to direct to a remote computer.
The Log Type list only contains logs that support syslog.
6. In the Destination field, type 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 displays 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

You can select a different log file or change the target computer for a remote log:

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Hardware.
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.
The Edit Remote Log window opens.
6. Change the log type or destination and click OK.
7. Be sure that the remote computer has the proper syslog configuration.

Deleting remote logs

Delete a remote log when it is no longer used.

1. Select a storage node in the navigation window and log in.
2. Open the tree below the storage node and select Hardware.
3. Select the Log Files tab.
4. Click Log File Tasks and select Delete Remote Log Destination.
A confirmation message opens.
5. Click OK.

NOTE:

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

9 Working with management groups

A management group is a collection of one or more storage nodes. It is the container within which you cluster storage nodes 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 nodes 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 node is not in a management group, the management group cannot use that storage node as a storage resource. For example, all of the storage nodes in a management group can be pooled into clusters for use by volumes in that group. To prevent a new storage node 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 nodes acts as the managers that control data transfer and replication.

Requirements for creating a management group

- IP addresses for the storage nodes that go into the group.
- Type of cluster you are planning: standard or Multi-Site.
- If a Multi-Site configuration, the physical sites and the storage nodes that go in them already created.
- Virtual IP addresses and subnet masks for the cluster.
- [Optional] Storage requirements for the volumes.

Managers overview

Within a management group, managers are storage nodes that govern the activity of all of the storage nodes in the group. All storage nodes contain the management software, but you must designate which storage nodes run that software by starting managers on them. These storage nodes then “run” managers, much like 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 nodes in the cluster.
- Re-synchronize data when storage nodes change states.
- Coordinate reconfigurations as storage nodes are brought up and taken offline.

One storage node has the coordinating manager. You can determine which storage node 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 node 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 is known as a “split-brain,” and 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 38](#).

Table 38 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 nodes 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 [Chapter 10](#) on page 189.

Failover managers

The Failover Manager is used in 2-node 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 node. 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 nodes in the SAN. [Figure 80](#) shows the Failover Manager installed, configured, and appearing in the CMC.

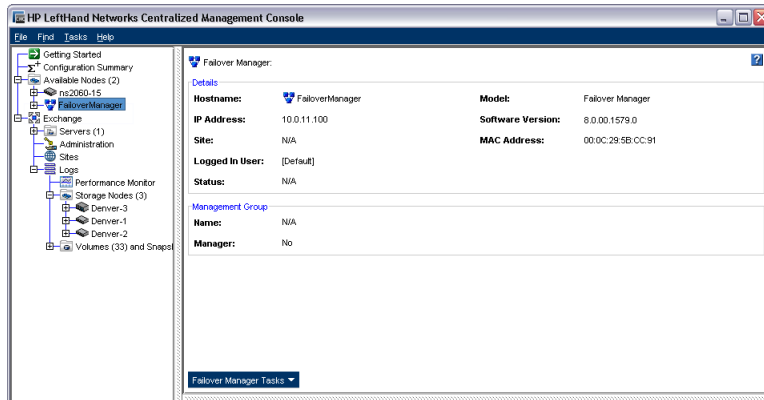


Figure 80 Failover manager in the available nodes pool

Once installed and configured, the Failover Manager operates as a storage node 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 81](#), but is not started on a storage node 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 node after quorum is lost. It is designed to be used in 2-node or 2-site system configurations which are at risk for a loss of quorum.

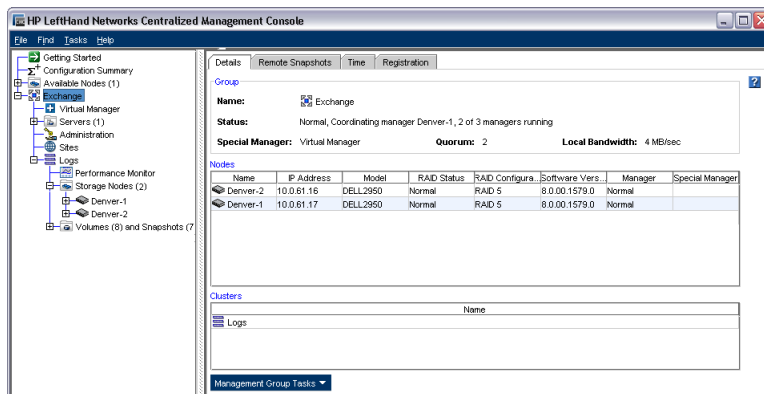


Figure 81 Virtual manager added to a management group

Creating a management group and default managers

When creating a management group, the wizard creates an optimal manager configuration for the number of storage nodes used to create the group. See [Table 39](#) 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.

Table 39 Default number of managers added when a management group is created

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

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 82](#). For each management group, the Configuration Summary displays an overview of the volumes, snapshots and storage nodes in that management group. The Summary roll-ups display configuration information and guides you to optimal configurations for volumes and snapshots, iSCSI sessions, and the number of storage nodes 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 nodes and iSCSI sessions contained in the management group.

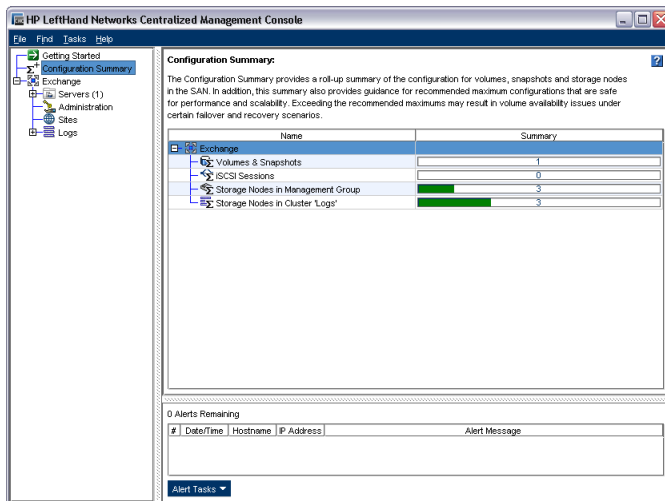


Figure 82 Configuration summary is created when the first management group is configured

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, then 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 displays orange for that line of the management group. Over 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 displays orange for that line of the management group. Over 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 nodes in the management group

The optimum number of storage nodes in a management group ranges up to 20. If the management group contains 21 to 30 storage nodes, the Configuration Summary displays orange for that line of the management group. Over 30 storage nodes triggers a warning by turning that line red. As soon as the total number of storage nodes reduces below the boundary, the summary bar returns to the previous indicator, either orange or green.

Storage nodes in the cluster

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

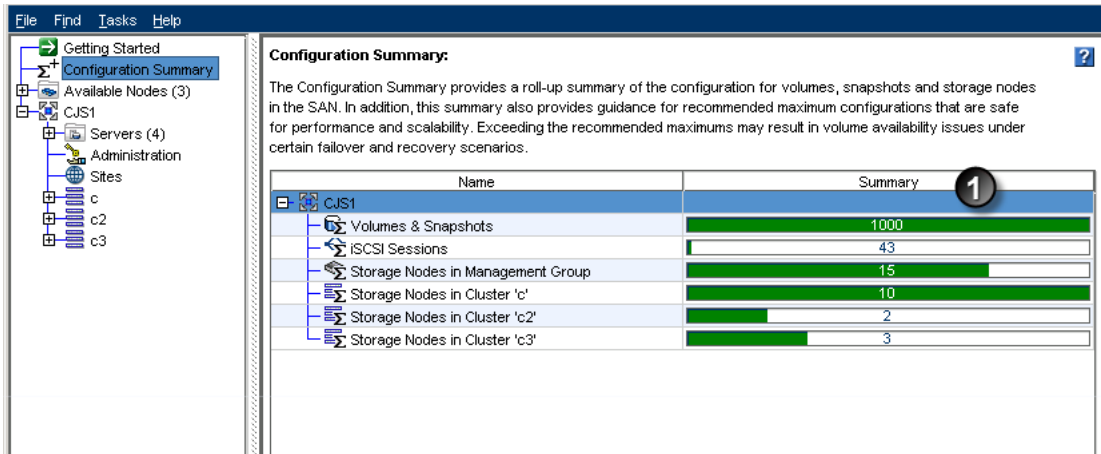
Reading the configuration summary

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 nodes, 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” on page 175.

Optimal configurations

Optimal configurations are indicated in green. For example, in Figure 83, there are 15 storage nodes in the management group “CJS1.” Those 15 storage nodes 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 nodes in cluster c3 are closer to the cluster recommended maximum for storage nodes than the 43 iSCSI sessions are to the maximum recommended iSCSI sessions for a management group.

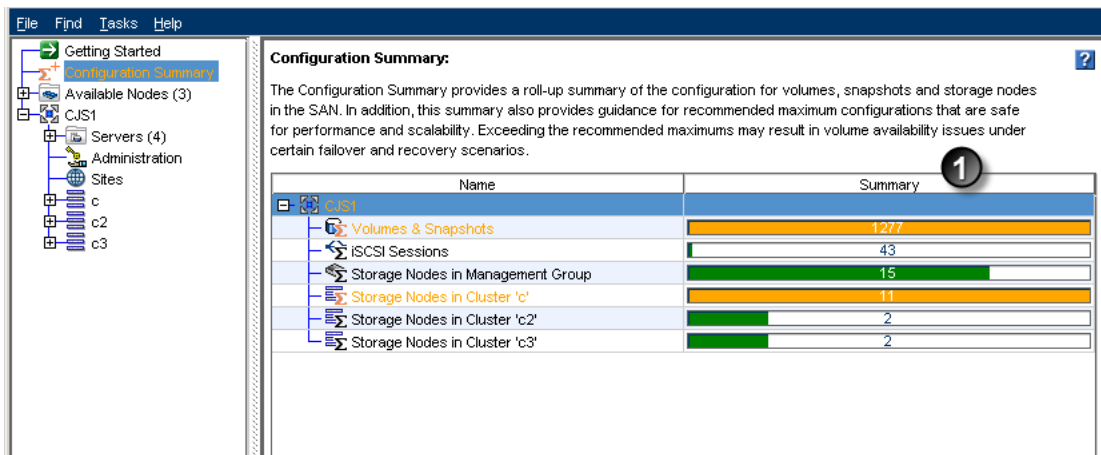


1. The items in the management group are all within optimum limits. The display is proportional to the optimum limits.

Figure 83 Understanding the summary graphs

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 84.

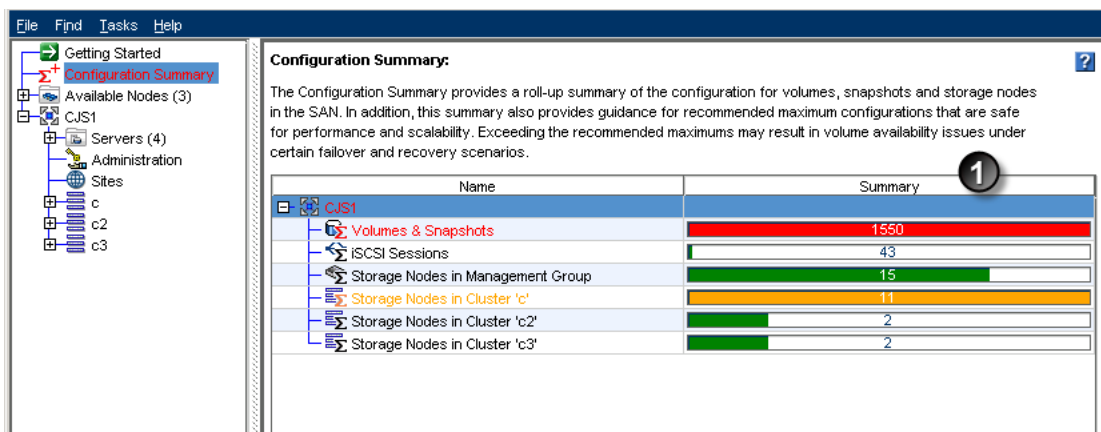


1. Volumes and snapshots are nearing the optimum limit. One cluster is nearing the optimum limit for storage nodes.

Figure 84 Warning when items in the management group are reaching safe limits

Configuration errors

When any item exceeds a recommended maximum, it turns red, and remains red until the number is reduced. See [Figure 85](#).



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

Figure 85 Error when some item in the management group has reached its limit

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

Guide to creating management groups

When using the Management Groups, Clusters and Volumes wizard, you must configure the characteristics described in [Table 40](#).

Table 40 Management group requirements

Management group requirement	What it means
Configure storage nodes	<p>Before you create a management group, make sure the storage nodes for the cluster are configured for monitoring, alerts and network bonding as best fits your network environment.</p> <p>CAUTION:</p> <p>[VSA] You cannot clone a VSA after it is in a management group. You must clone a VSA while it is in the Available Nodes pool.</p>
Plan administrative users	<p>When you create a management group, you must add the first administrative user. This user has full administrative permissions. Add additional users later. See Adding a new administrative user (page 123).</p>
Plan date and time configuration	<p>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 Chapter 5.</p>
Plan Virtual IP Addresses (VIPs)	<p>A VIP is required for each cluster. VIPs ensure fault tolerant server access to the cluster and enable iSCSI load balancing. See Configure virtual IP and iSNS for iSCSI (page 210).</p>
Starting a manager	<p>A management group must have an optimum number of managers running. The Management Groups, Clusters and Volumes wizard attempts to set the proper number of managers for the number of storage nodes you use to create the management group.</p>
Assigning manager IP addresses	<p>The storage nodes that are running managers must have static IP addresses (or reserved IP addresses if using DHCP). That is, the IP address must not change while that storage node is a manager.</p>

Getting there

You create a management group using the Management Groups, Clusters and Volumes wizard. Access the wizard in any of the following ways:

- From the Getting Started Launch Pad, by selecting the Management Groups, Clusters and Volumes wizard. See [“Creating storage by using the Getting Started Launch Pad”](#) on page 37.
- By right-clicking an available storage node in the navigation window.
- From the menu bar with Tasks > Management Group > New Management Group.

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.

Create management group and add storage nodes

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

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 nodes 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 nodes in this management group.
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 LeftHand P4000 Multi-Site HA/DR Solution Pack User Manual.

1. Select Standard Cluster, then click Next.
2. Type a cluster name in the Create a Cluster window.
3. From the list select the storage nodes to include in the cluster.
4. Click Next to assign Virtual IPs.
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, replication level, size and provisioning type for the volume.
2. Click Finish.

After a few minutes, a summary window opens, listing details of the new management group, cluster and volume.

3. Click Close.
4. A message opens notifying you that you must register.
This registration is required to use advanced features such as multi-node clusters and Remote Copy. For more information about registering advanced features, see [Chapter 19](#) on page 317.
5. Click OK.
The navigation window displays the new management group, cluster with storage nodes, and volumes.
6. As a last step, back up the configuration data of the entire management group.
See [“Backing up a management group configuration”](#) on page 183.

Adding a storage node to an existing management group

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

1. In the navigation window, select an available storage node that you want to add to a management group.
2. Click Storage Node 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 node to run a manager, select the storage node in the management group, right-click, and select Start Manager.
6. Repeat [Step 1](#) through [Step 4](#) to add additional storage nodes.
7. Save the configuration data of the changed management group.
See [“Backing up a management group configuration”](#) on page 183.

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 node in a management group, you are logged in to all storage nodes in that group.

Choosing which storage node to log in to

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

1. When the Log in to Node window opens, click Cancel.
A message opens, asking if you want to log in to a different storage node.
2. Click OK.
3. The Log in to Node window opens with a different storage node listed.
4. If that is the storage node you want, go ahead and log in. If you want to log in to a different storage node, repeat [Step 1](#) and [Step 2](#) until you see the storage node you want.

Logging out of a management group

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

Management group Maintenance tasks

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

- [“Starting and stopping managers”](#) on page 181
- [“Editing a management group”](#) on page 182
- [“Backing up a management group configuration”](#) on page 183
- [“Restoring a management group”](#) on page 184
- [“Safely shutting down a management group”](#) on page 184
- [“Start the management group back up”](#) on page 185
- [“Removing a storage node from a management group”](#) on page 187
- [“Deleting a management group”](#) on page 187

Starting and stopping managers

After adding the storage nodes to the management group, start managers on the additional storage nodes in the management group. The number of managers you start depends upon the overall design of your storage system. See [“Managers overview”](#) on page 171 for more information about how many managers to add.

Starting additional managers

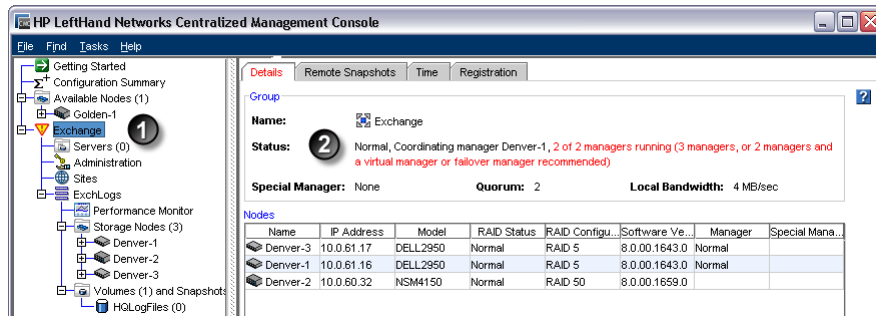
1. In the navigation window select a storage node in the management group on which to start a manager.
2. Click Storage Node Tasks on the Details tab and select Start Manager.

Repeat these steps to start managers on additional storage nodes.

Stopping managers

Under normal circumstances, you stop a manager when you are removing a storage node from a management group. You cannot stop the last manager in a management group. If you stop a manager

that compromises fault tolerance, the management group displays the icon indicating an item needs attention.



1. Management group displays icon
2. Status describes quorum risk

Figure 86 Manager quorum at risk

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

Implications of stopping managers

- Quorum of the storage nodes may be decreased.
- Fewer copies of configuration data is 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 node for which you want to stop the manager.
3. Click Storage Node 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 includes the following items:

- Changing local bandwidth priority.
- Editing Remote Bandwidth, done on the management group containing the remote snapshot. (See the Remote Copy User Manual, 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 “[Backing up a management group configuration](#)” on page 183.

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 .25 MB/sec.

Local bandwidth priority settings

The bandwidth setting is in MB per second. Use [Table 41](#) as a guide for setting the local bandwidth.

Table 41 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 displays on the Details tab in the management group tab window.

Backing up a management group configuration

Use Backup Configuration of Management Group to save one or both of the following configuration files:

- Back up the configuration—creates a binary file (.bin) of the management group configuration

- Save the configuration description—creates a text file (.txt) listing the configuration characteristics of the management group

The binary file enables you to automatically recreate a management group with the same configuration. Use the text file for support information. Your support representative will help you restore this back up.

 **NOTE:**

Backing up the management group configuration does not save the configuration information for the individual storage nodes in that management group nor the data. To back up storage node configurations, “[Backing up the storage node configuration file](#)” on page 46.

Backing up a management group with remote copy relationships

If you back up a management group that is participating in Remote Copy, it is important to back up the associated Remote Copy management groups at the same time. If you back them up at different times, and then try to restore one of the groups, the back up files will not match. This mismatch will cause problems with the restore.

Backup a management group configuration

1. In the navigation window, select the management group and log in.
2. Click Management Group Tasks on the Details tab and select View Management Group Configuration.
3. Click Save.
A Save window opens so that you can select the location for the .bin file or .txt file.
4. In the Save window, accept the default name of the file, or type a different name.
5. From the Files of Type drop-down menu, select the .bin file type.
6. Repeat this procedure, and in Step 5, select the .txt file type.

The .txt file describes the configuration.

Restoring a management group

Call Customer Support for help if you need to restore a management group using a .bin 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 nodes in that group
- Move storage nodes 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 [Chapter 16](#). Sample scripts are available from the Customer Resource Center.

Shutting down a management group also relates to powering off individual storage nodes and maintaining access to volumes. See the command line documentation, the *Cliq User Manual*, 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.

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.

The management group shuts down and disappears from the CMC.

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.

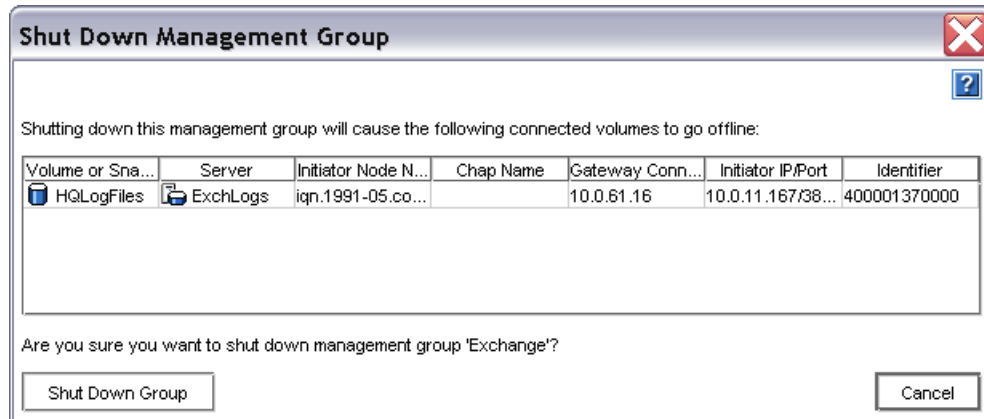


Figure 87 Notification of taking volumes offline

1. Stop server or host access to the volumes in the list.
2. Click Shut Down Group.

The management group shuts down and disappears from the CMC.

Start the management group back up

When you are ready to start up the management group, simply power on the storage nodes for that group.

1. Power on the storage nodes that were shut down.
2. Use Find in the CMC to discover the storage nodes.

When the storage nodes 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 node becomes unavailable, and the management group is shut down while that storage node is repaired or replaced. After the storage node is repaired or replaced and the management group is started up, the management group remains in maintenance mode while the repaired or replaced storage node is resynchronizing with the rest of the management group.
- After a management group is shut down, a subset of storage nodes is powered on. The management group remains in maintenance mode until the remaining storage nodes are powered on and re-discovered in the CMC.
- For some reason, a storage node 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 nodes 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.

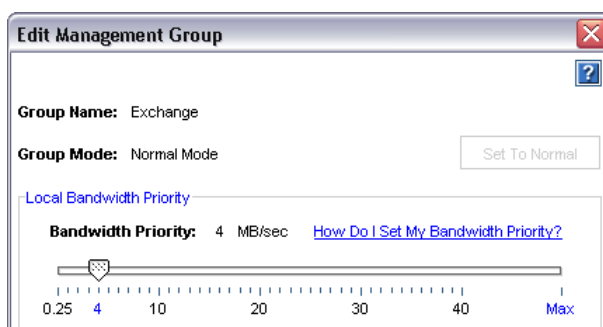


Figure 88 Manually setting management group to normal mode

3. Click Set To Normal.
The management group is reset to normal mode.

Removing a storage node from a management group

Prerequisites

- Stop the manager on the storage node if it is running a manager. You may want to start a manager or Virtual Manager on a different storage node to maintain quorum and the best fault tolerance. See “[Stopping managers](#)” on page 181.
- (Optional) If the resultant number of storage nodes in the cluster is fewer than the volume replication level, you may have to reduce the replication level of the volume(s) on the cluster before removing the storage node from the cluster.
- Remove the storage node from the cluster. See “[Removing a storage node from a cluster](#)” on page 212.
- Let any restripe operations finish completely.

Remove the storage node

1. Log in to the management group from which you want to remove a storage node.
2. In the navigation window, select the storage node to remove.
3. Click Storage Node Tasks on the Details tab and select Remove from Management Group.
4. Click OK on the confirmation message.

In the navigation window, the storage node is removed from the management group and moved to Available Nodes 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 nodes 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 nodes return to the Available Nodes 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.

10 Using specialized managers

The SAN/iQ software provides two specialized managers that are used in specific situations. The Failover Manager is used in 2-node 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 not started on a storage node until a failure in the system causes a loss of quorum. It is designed to be used in 2-node or 2-site system configurations which are at risk for a loss of quorum.

Definitions

Terms used in this chapter:

- **Virtual manager**—A manager which is added to a management group but is not started on a storage node until a failure in the system causes a loss of quorum. The virtual manager is designed to be used in specific system configurations which are at risk of a loss of quorum.
- **Failover Manager**—A specialized manager running as a VMware guest operating system that can act as a quorum tie-breaker node when installed into a 3rd location in the network to provide for automated failover/failback of the Multi-Site SAN clusters.
- **Regular manager**—A manager which is started on a storage node and operates according to the description of managers found in “[Managers overview](#)” on page 171.
- **Manager**—any of these managers.

Failover manager overview

The Failover Manager is a specialized version of the SAN/iQ software designed to run as a virtual appliance in a VMware environment. The Failover Manager participates in the management group as a real manager in the system; however, it does 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 physical hardware in a site.

Failover Manager requirements

- Static IP address, or a reserved IP address if using DHCP.
- Bridged connection through the VMware Console, or assigned network in VI Client.
- Server on which to install the Failover Manager.
- Only one Failover Manager per management group.
- You cannot have a virtual manager and a Failover Manager in the same management group.
- You cannot run the Failover Manager inside a virtual Windows machine with VMware Server running.

Minimum system requirements for using with VMware server or player

- 10/100 ethernet
- 384 MB of RAM

- 5 GB of available disk space
- VMware Server 1.x

Minimum system requirements for using with VMware ESX Server

- VMware ESX Server version 3.x
- 1024 MB of RAM

△ CAUTION:

Do not install the Failover Manager on the HP LeftHand Storage Solution, since this would defeat the purpose of the Failover Manager.

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.
- Failover Manager directory, host name and IP address.
- 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.

Upgrading the 7.0 Failover Manager

The Failover Manager released with SAN/iQ software version 7.0 cannot be upgraded or patched. To upgrade to the Failover Manager released with SAN/iQ software version 8.0 you must uninstall the previous version. From version 8.0 on you will be able to upgrade the Failover Manager.

1. Remove the Failover Manager from the management group.
2. Uninstall the Failover Manager, as described in “[Uninstalling the Failover Manager for VMware Server or Player](#)” on page 194.
3. Install the new version of the Failover Manager.
4. Configure the name and IP address.
5. Add the new Failover Manager to the management group.

Using Failover Manager with VMware Server or VMware Player

Installing and configuring the Failover Manager

Install the Failover Manager from the HP LeftHand VSA CD or download it from the HP LeftHand Networks web site.

Failover Manager configuration

When the Failover Manager is installed, it is automatically configured as follows:

- To auto-start in the event either the VMware Console or host server reboots.
- The virtual network adapters under the Failover Manager are configured as Bridged network adapters.

After the Failover Manager is installed, it is started in the VMware Console. After it is started, you use the Configuration Interface to set the IP address.

To install the Failover Manager

Install the Failover Manager onto a separate server on the network.

△ CAUTION:

Do not install the Failover Manager on the HP LeftHand Storage Solution, since this would defeat the purpose of the Failover Manager.

Using the HP LeftHand Management DVD

1. Using the HP LeftHand Management DVD, click Install on the opening window.
The install software window opens.
2. Click Failover Manager.
3. Continue through the installation wizard, following the instructions on each window.
After the installation wizard finishes, the default choice is to launch the Failover Manager.
4. Click Finish to exit the wizard and start the Failover Manager.

Using the HP LeftHand Networks web site download

1. Click download on the web site, and the installation wizard opens.
2. Continue through the installation wizard, following the instructions on each window.
After the installation wizard finishes, the default choice is to launch the Failover Manager.
3. Click Finish to exit the wizard and start the Failover Manager.

To configure the Failover Manager

The system pauses while the Failover Manager is installed, registered and then the VMware Console opens, as shown in [Figure 89](#).

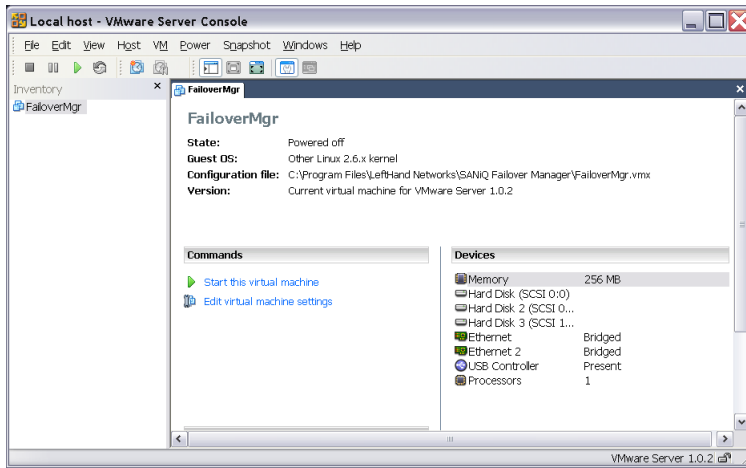


Figure 89 VMware Console opens with Failover Manager installed and registered

The Failover Manager then automatically powers on.

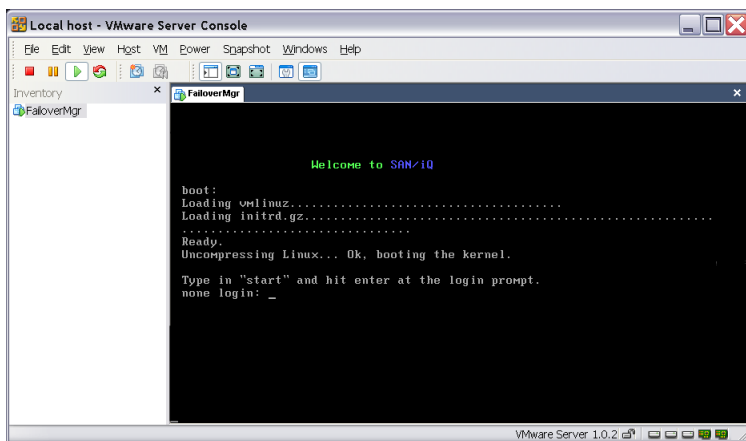


Figure 90 Failover Manager boots up

1. At the system login prompt, click in the window, type start and press Enter.
To get the cursor back from the VMware Console, press Ctrl+Alt.
The Configuration Interface Login opens.
2. Press Enter to open the Configuration Interface main menu.
3. Tab to Network TCP/IP Settings and press Enter.
The Available Devices window opens.

4. Make sure eth0 is selected and press Enter.

The Network Settings window opens.

```
+-[ Network Settings: ]-----+
|
| Specify the network settings for the unknown port. Be
| sure the ethernet cable is plugged into the selected
| port.
|
| Hostname:  FailoverManager_
|
| ( ) Disable Interface.
| ( ) Obtain IP address automatically using DHCP.
| (*) Use the following IP address:
|
|   IP Address:  10.0.14.87
|   Mask:        255.255.0.0
|   Gateway:     0.0.0.0
|
|                   [ OK ] [ CANCEL ]
|
+-----+
```

Figure 91 Setting the host name and IP address

5. Tab to the choice that reflects how you want to set the IP address for the Failover Manager. If you use DHCP to assign the IP address, be sure to reserve that IP.
6. Tab to OK and press Enter. The Modify Network Settings confirmation window opens.
7. Tab to OK and press Enter to confirm the change. After a minute or less, the IP address window opens.

```
+-[ Available Network Devices: ]+
|
| [ eth0 ]
|
+-----+
|
| Specify the network settings for the unknown port. Be
|
+-----+
|
| The IP Address of this Storage Server has been set to:
| 10.0.14.87
|
|
|                   < OK >_
|
+-----+
```

Figure 92 Confirming the new IP address

8. Make a note of the IP address and press Enter to close the IP address window. The Available Network Devices window opens.
9. Tab to Back and press Enter. The Configuration Interface main menu opens.
10. Tab to Log Out and press Enter. The Configuration Interface login window opens.
11. Click File > Exit to close the VMware Console.

Uninstalling the Failover Manager for VMware Server or Player

Use the SAN/iQ Management Software DVD to uninstall the Failover Manager.

1. Remove the Failover Manager from the management group.
2. Insert the SAN/iQ Management Software DVD into the CD/DVD drive and click Install on the opening window.
3. Click Failover Manager.
4. The Failover Manager installation wizard opens.
5. Click through the wizard, selecting Uninstall from the Repair or Uninstall window when it opens. The Failover Manager is removed from the server.

Troubleshooting the Failover Manager on VMware Server or Player

Two issues may occur when running a Failover Manager.

- The Failover Manager may not automatically restart upon a reboot. The default Startup/Shutdown Options settings may have accidentally changed.
- You cannot find the Failover Manager in the CMC, nor ping it on the network. The default configuration may bridge to an incorrect host adapter.

Use the following instructions to correct the VMware Server settings for either of these issues.

Fix startup/shutdown options

1. Open the VMware Server Console.
2. Select the Failover Manager virtual machine in the Inventory list.
3. Power off the Failover Manager.
4. From the menu, select VM > Settings, or right-click on the Failover Manager and select Settings. The Virtual Machine Settings window opens.
5. Select the Options tab, and select Startup/Shutdown.
6. On the right side in the Virtual machine account section, select an account for the virtual machine, such as Local system account.
7. On the right side in the Startup/Shutdown Options section, select the following options:
 - On host startup:
Power on virtual machine
 - On host shutdown:
Power off virtual machine
8. When you are done, click OK to save your changes.
9. Power on the Failover Manager.

Fix network settings to find Failover Manager

1. Determine which interface on the windows host server is configured for the SAN networks.
2. Open the VMware Server Console.
3. Select the Failover Manager virtual machine in the Inventory list.
4. From the menu, select Host > Virtual Network Settings.
5. Click the Automatic Bridging tab.
6. Ensure that the checkbox under Automatic Bridging is checked.
7. Click Add in the Excluded Adapters section.
A list of network adapters opens.
8. Add all adapters in the list except the adapter identified in step 1.
9. Click OK when you are finished.

Using the Failover Manager on VMware ESX Server

CAUTION:

Do not install the Failover Manager on the HP LeftHand Storage Solution, since this would defeat the purpose of the Failover Manager.

Installing the Failover Manager on VMware ESX Server

Install the Failover Manager from the HP LeftHand Management DVD or obtain the Failover Manager as a zip package downloaded from the HP LeftHand Networks web site.

When you install the Failover Manager for the first time, you will need to

- Start the VMware Infrastructure Client (VI Client).
- Transfer or upload the virtual machine to the ESX Server.
- Add the Failover Manager to inventory.
- Power on the Failover Manager.
- Set the IP address and host name of the Failover Manager.

NOTE:

By default, ssh and scp commands are disabled for the root user. To enable access, see the VMware documentation for ESX Server Basic Administration.

Using the HP LeftHand Management DVD

1. Using the HP LeftHand Management DVD, click Install on the opening window.
The install software window opens.
2. Click Failover Manager for ESX.

3. Continue through the installation wizard, following the instructions on each window.

Using the HP LeftHand Networks web site download

1. Click download on the web site, and the installation wizard opens.
2. Continue through the installation wizard, following the instructions on each window.
After the installation wizard finishes, the default choice is to launch the Failover Manager.
3. Click Finish.

Installing the Failover Manager Files on the ESX Server

Use one of the following methods, depending on your software.

For ESX 3.5+ or ESXi

1. Connect to ESXi host via VC or VI Client.
2. Click on ESXi host and go to the Configuration tab.
3. Select Storage.
4. Find the local VMFS datastore where the Failover Manager will be hosted.
5. Right-click and select Browse DataStore.
6. Create a new directory and click on the upload files icon.

For ESX Server 3.0 to 3.0.2

1. Upload unzipped folder for the Failover Manager.
2. Make a directory for the Failover Manager at
`/vmfs/volumes/"your_datastore_name"`.
3. Copy the Failover Manager files into the directory you just created on the ESX Server using `scp` (Linux) or `pscp` (Windows), as shown in the example below.

```
scp *.* <user>@<IP address of the ESX Server>:/vmfs/volumes/datastore
```
4. Open execute permissions on the `.vmx` file using the command `chmod 755 FOM.vmx`.

Configuring the Failover Manager using the VI Client

After you have installed the Failover Manager files on the ESX Server, you configure the Failover Manager using the VI Client.

Add Failover Manager to inventory

1. In the Inventory Panel, select the VMware ESX Server.
2. In the Information Panel, select the Configuration Tab.
3. In the Hardware section, select Storage (SCSI, SAN, and NFS).
4. In the Storage section, right-click the datastore icon and select Browse Datastore.
The Datastore Browser opens.

5. Right-click the FailoverMgr.vmx file and select Add to Inventory.
6. In the Add to Inventory Wizard, enter a name for the new Failover Manager and click Next.
7. Select the Inventory Locations to place the Failover Manager in the Add to Inventory wizard.
8. Verify the information and click Finish.
9. Close the DataStore Browser.

Select a network connection

1. In the Inventory Panel select the Failover Manager.
2. In the Information Panel select the Summary tab. In the Commands section, select Edit Settings. The Virtual Machine Properties window opens.
3. On the Hardware tab, select Network Adapter 1.
4. Select the appropriate network connection from the Network label list on the right.
5. Click OK to exit the Virtual Machine Properties window.

Power on the Failover Manager and configure the IP address and host name

1. In the Inventory panel, select the new Failover Manager and power it on using the Power On command on the Information panel.
2. Click the Console Tab and wait for the Failover Manager to boot.
3. When the Failover Manager has finished booting, a log in prompt opens.

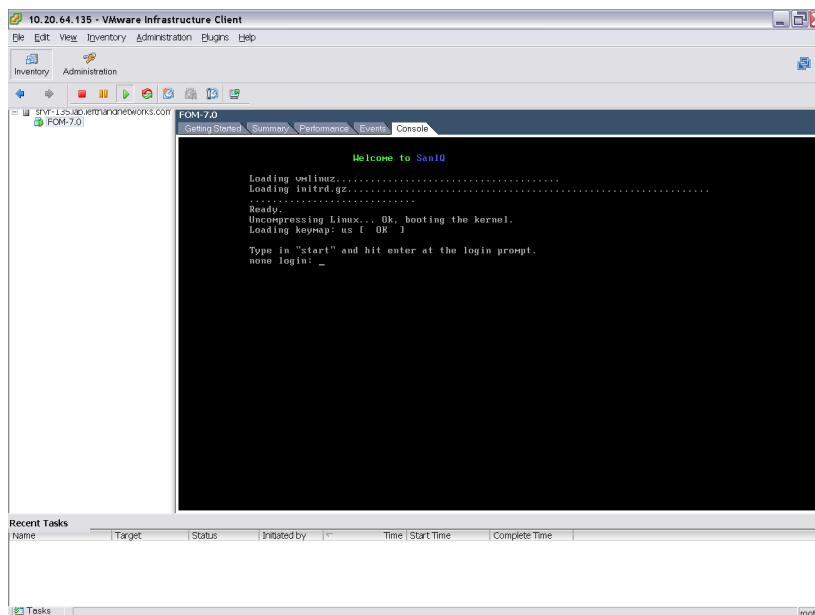


Figure 93 Logging in to the SAN/iQ Configuration Interface

4. Log in and use the SAN/iQ Configuration Interface to configure an IP address and host name for the Failover Manager.

Setting the IP address

Use the Configuration Interface to set the IP address for the Failover Manager.

1. Enter start and press Enter.
2. Press Enter to log in.
3. Tab to Network TCP/IP Settings and press Enter.
4. Tab to the network interface and press Enter.
5. Tab to the Hostname field if necessary.
6. Press Backspace in the Hostname field to delete the default name and enter your own host name.
This host name displays in the CMC. It does not change the name of the original FOM.vmx file, nor does it change the name in VMware.
7. Configure the IP Address one of two ways.

Using DHCP	Configure IP address manually
<ol style="list-style-type: none"> 1. Tab to the choice Obtain IP address automatically using DHCP and press Enter to select it. 2. Tab to OK and press Enter. A message opens, asking you to verify the request. 3. Tab to OK and press Enter. After a short pause, another message opens that displays the new IP address. Record this IP address for later use. 	<ol style="list-style-type: none"> 1. Tab to the choice Use the Following IP Address and press Enter. The IP Address, Netmask, and Gateway list opens for editing. 2. Tab to each field and enter the appropriate information. Gateway is a required field. If you do not have a gateway, enter 0.0.0.0.

8. Tab to OK and press Enter.
A confirmation message opens.
9. Press Enter.
The network interface is now configured (wait a few seconds). The Available Network Devices window opens.
10. On the Available Network Devices window, tab to Back and press Enter to return to the Configuration Interface menu.
11. Tab to Log Out and press Enter.
The Configuration Interface entry window is displayed again.
12. Press Ctrl+Alt to get the cursor back from the Console.

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” then they are still running correctly. The “out of date” status is not a problem. VMware tools are updated with each SAN/iQ software upgrade.

3. In the Inventory panel, right-click the Failover Manager and select Rename.

4. Change the name of the Failover Manager to match the host name, if desired.
Your Failover Manager is ready to use.
5. Minimize your VI Client session.

Next, use the Find function to discover the Failover Manager in the CMC and then add the Failover Manager to a management group.

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.

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 42 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. 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 node if it can be found. • Open a VI Client session and select the Summary tab for the node 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, then press any key.

Virtual manager overview

A virtual manager is a manager that is added to a management group, but is not started on a storage node until it is needed to regain quorum. A virtual manager provides disaster recovery for one of two configurations:

- Configurations with only 2 storage nodes. (A virtual manager will automatically be added when creating a management group using 2 storage nodes.)
- Configurations in which a management group spans 2 geographic sites.

See “[Managers and quorum](#)” on page 172 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 node goes offline, it can also be used for maintaining quorum during maintenance procedures.

When to use a virtual manager

Use a virtual manager in the following configurations:

- A management group across two sites with shared data
- A management group in a single location with two storage nodes

Use a virtual manager for disaster recovery in a two-site configuration, or a two-node configuration. You can also use a virtual manager to maintain quorum during storage node 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 nodes in both locations.

Management group in a single location with two storage nodes

If you create a management group with only two storage nodes, that management group is not a fault tolerant configuration. Using one manager provides no fault tolerance. Using two managers also provides no fault tolerance, due to loss of quorum if one manager becomes unavailable. See “[Managers and quorum](#)” on page 172 for more information.

Running two managers and adding a virtual manager to this management group provides the capability of regaining quorum if one manager becomes unavailable.

Storage node 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 node offline for maintenance.

Benefits of a virtual manager

Running a virtual manager supports disaster tolerant configurations to support full site failover. The virtual manager ensures that, in the event of either a failure of a storage node running a manager, or of communication breakdown between managers (as described in the two-site scenario), quorum can be recovered and, hence, data remains accessible.

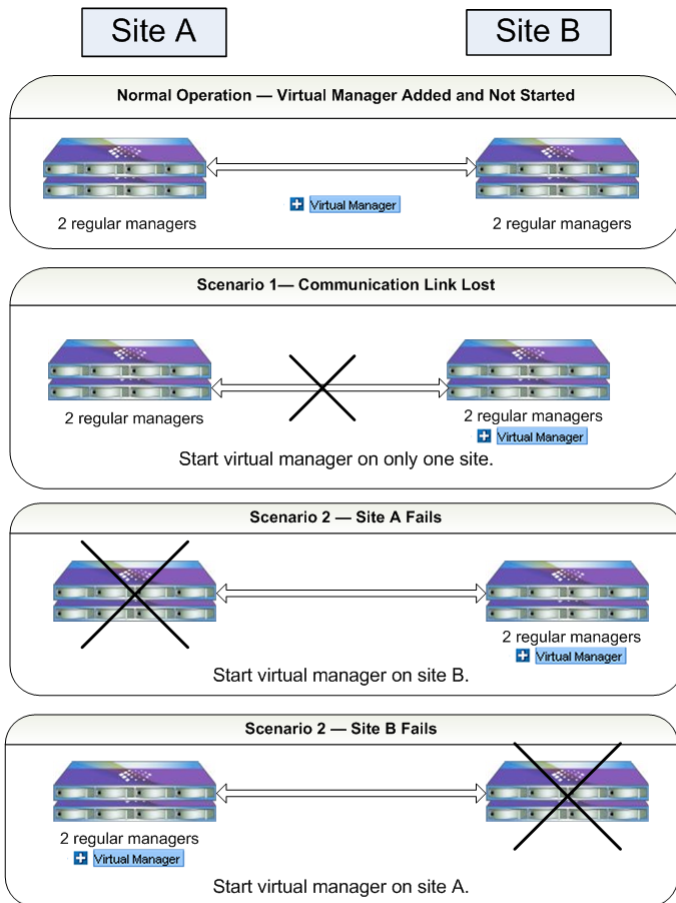
Requirements for using a virtual manager

It is critical to use a virtual manager correctly. A virtual manager is added to the management group, but not started on a storage node until the management group experiences a failure and a loss of quorum. To regain quorum, you start the virtual manager on a storage node that is operating and in the site that is operational or primary.

Table 43 Requirements for using a virtual manager

Requirement	What it means		
Use a virtual manager with an even number of regular managers running on storage nodes	Disaster Recovery Scenario	Number of regular managers running	Total number of managers, including the virtual manager
	2 separate sites with shared data	4	5
	2 storage nodes 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.		

Illustrations of correct uses of a virtual manager are shown in [Figure 94](#).



- Examples of 2-site failure scenarios where a virtual manager is started to regain quorum.
- In all the failure scenarios, only one site becomes primary with a virtual manager started.

Figure 94 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 nodes in one cluster. The cluster spans two geographic sites with two storage nodes at each site. The cluster contains a single volume with 2-way replication 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 nodes with site-identifying host names.

To ensure that you can easily identify which storage nodes reside at each site, use host names that identify the storage node location. See [“Changing the storage node hostname”](#) on page 44 .

Management Group Name—TransactionData

Storage Node 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 2-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”](#) on page 171.

3. Add storage nodes to the cluster in alternating order.

Create the cluster. Add the storage nodes to the cluster in alternating order of site , as shown in the bulleted list. The order in which the storage nodes are added to the cluster determines the order in which copies of data is written to the volume. Alternating the addition of storage nodes by site location ensures that data is written to each site as part of the 2-way replication you configure when you create the volume. See [“Creating additional clusters”](#) on page 209.

Cluster Name—CreditData

Add storage nodes to cluster in the following order

- 1st storage node—Boulder-1
- 2nd storage node—Golden-1
- 3rd storage node—Boulder-2
- 4th storage node—Golden-2

△ CAUTION:

If storage nodes are added to the cluster in any order other than alternating order by site, you will not have a complete copy of data on each site.

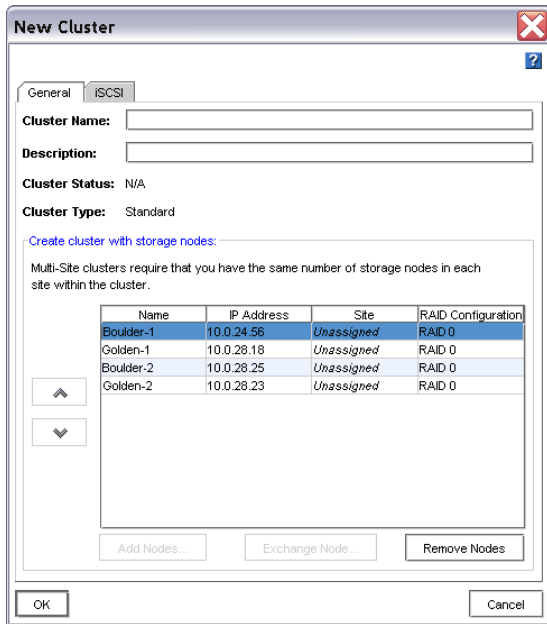


Figure 95 Adding storage nodes to cluster in alternating site order

4. Create the volume with 2-way replication.

Two-way replication causes two copies of the data to be written to the volume. Because you added the storage nodes to the cluster in alternating order, a complete copy of the data exists on each site. See “Planning data replication” on page 223.

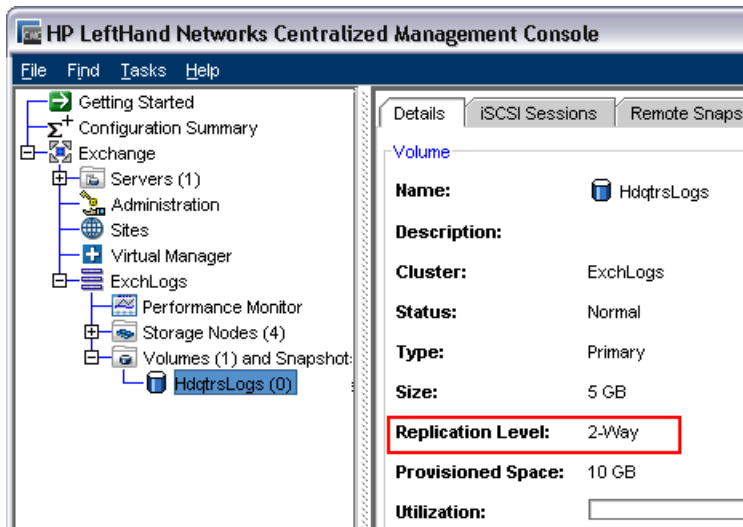


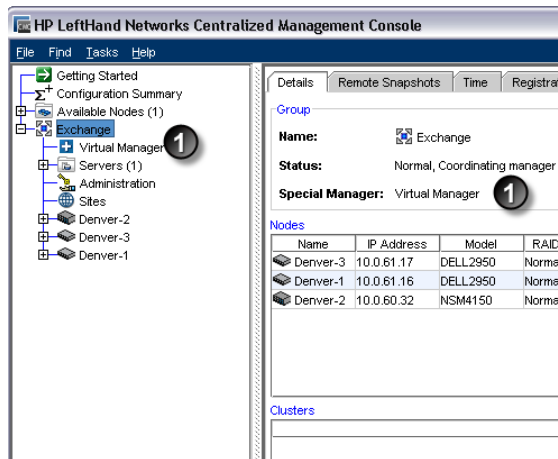
Figure 96 2-way replicated volume on 2-site cluster

Adding a virtual manager

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.
A confirmation message opens.

3. Click OK to continue.

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.



1. Virtual manager added

Figure 97 Management group with virtual manager added

The virtual manager remains added to the management group until needed.

Starting a virtual manager to regain quorum

Only start a virtual manager when it is needed to regain quorum in a management group. [Figure 94](#) (page 202) 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 2-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 nodes 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 reestablish communication and ensure that the data in both sites are 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 nodes and reconstruct the cluster. Refer to customer support for help with cluster recovery. You must have the serial number of one of your storage nodes when making a support call.

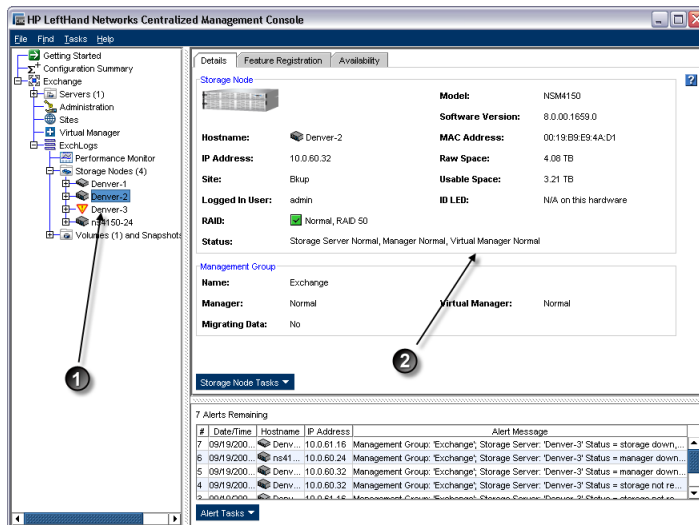
Two-site scenario, communication between the sites is lost

In this scenario, the sites are both operating independently. On the appropriate site, depending upon your configuration, select one of the storage nodes 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 are resynchronized. When the data is resynchronized, stop the virtual manager to return to the disaster tolerant configuration.

Starting a virtual manager

A virtual manager must be started on a storage node, ideally one that isn't already running a manager. However, if necessary, you can start a virtual manager on a storage node that is already running a manager.

1. Select the storage node on which you want to start the virtual manager.
2. Click Storage Node Tasks on the Details tab and select Start Virtual Manager.



1. Unavailable Manager
2. Virtual manager started

Figure 98 Starting a virtual manager when storage node running a manager becomes unavailable

NOTE:

If you attempt to start a virtual manager on a storage node that appears to be normal in the CMC, and you receive a message that the storage node is unavailable, start the virtual manager on a different storage node. This situation can occur when quorum is lost because the CMC may still display the storage node in a normal state, even though the storage node is unavailable.

Verifying virtual manager status

Verify whether a virtual manager has been started, and if so, which storage node 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 node with the virtual manager.
2. Click Storage Node Tasks on the Details tab and select Stop Virtual Manager.

A confirmation message opens.

3. Click OK.

The virtual manager is stopped. However, it remains part of the management group and part of the quorum.

Removing a virtual manager

You can remove the virtual manager from the management group altogether.

1. Select the management group from which you want to remove the virtual manager and log in.
2. Click Management Group Tasks on the Details tab and select Delete Virtual Manager.

A confirmation window opens.

3. Click OK to continue.

The virtual manager is removed.

NOTE:

The CMC will not allow you to delete a manager or virtual manager if that deletion causes a loss of quorum.

11 Working with clusters

Within a management group, you create subgroups of storage nodes called clusters. A cluster is a grouping of storage nodes from which you create volumes. Volumes seamlessly span the storage nodes in the cluster.

Think of a cluster as a pool of storage. You add storage to the pool by adding storage nodes. You then carve volumes and snapshots out of the pool.

Before creating a cluster, make sure you are familiar with the iSCSI information in [Chapter 22](#) on page 335.

Clusters and storage node capacity

Clusters can contain storage nodes with different capacities. However, all storage nodes in a cluster operate at a capacity equal to that of the smallest capacity storage node.

Prerequisites

- Before you create a cluster, you must have created a management group.

Creating additional clusters

When you create a management group, you create the first cluster in that management group. Use the following steps to create additional clusters in existing management groups.

Prerequisites

- An existing management group
- At least one storage node in the management group that is not already in a cluster.

Number of storage nodes in clusters

For information about the recommended maximum number of storage nodes that can be added safely to a cluster, see [“Configuration summary overview”](#) on page 174 or [Chapter 9](#) on page 171.

To create additional clusters

1. Log in to the management group for which you want to create a cluster.
2. Right-click on the storage node and select Add to Existing or New Cluster.
3. Select New Cluster and click Add.
4. Enter a meaningful name for the cluster.

A cluster name is case sensitive and must be from 1 to 127 characters. It cannot be changed after the cluster is created.

5. (Optional) Enter a description of the cluster.
6. Select one or more storage nodes from the list.
Use the up and down arrows on the left to promote and demote storage nodes in the list to set the logical order in which they appear. For information about one specific disaster recovery configuration when the order matters, see [“Configuring a cluster for disaster recovery”](#) on page 202.
7. Click the iSCSI tab.

Configure virtual IP and iSNS for iSCSI

VIPs are required for iSCSI load balancing and fault tolerance and for using HP LeftHand DSM for MPIO. For more information, see [Chapter 22](#) on page 335.

New for release 8.0

Virtual IP (VIP) addresses are required for all clusters in SAN/iQ software versions 8.0 and higher.

1. Click the iSCSI tab to bring it to the front.
Because a VIP is required in release 8.0 and greater, the choice to use a virtual IP is disabled by default in the 8.0 CMC. If you have management groups that are running 7.0 or earlier software, the choice to use a VIP remains enabled.
2. Add the IP address and subnet mask.

Adding an iSNS server

(Optional) Add an iSNS server.

NOTE:

If you use an iSNS server, you may not need to add Target Portals in the Microsoft iSCSI Initiator.

1. In the iSCSI tab view, open the iSCSI Tasks menu and click Add iSNS Server.
The Add iSNS Server window opens.
2. Enter the IP address of the iSNS server.
3. Click OK.
4. Click OK when you have finished.
The cluster is created and displayed inside the management group.
5. Select the cluster to open the Clusters tab window.

Tracking cluster usage

The Use Summary, Volume Use and Node Use tabs provide detailed information about provisioning of volumes and snapshots and space usage in the cluster. See [“Ongoing capacity management”](#) on page 227 for information about the information reported on these tabs.

 **NOTE:**

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 associated with the cluster.

Editing a cluster

When editing a cluster, you can change the description, and add or remove storage nodes. You can also edit or remove the virtual IP and iSNS servers associated with the cluster.

Prerequisite

You must log in to the management group before you can edit any clusters within that group.

Getting there

1. In the navigation window, select the cluster you want to edit.
2. Click Cluster Tasks and select Edit Cluster.

Adding a new storage node to an existing cluster

Add a new storage node to an existing cluster to expand the storage for that cluster.

 **NOTE:**

Adding a storage node to a cluster causes a restripe of the data in that cluster. A restripe may take several hours or longer.

Adding a new storage node is not the same as replacing a repaired storage node with a new one. If you have repaired a storage node and want to replace it in the cluster, see [“Repairing a storage node”](#) on page 216.

Prerequisite

- Add the storage node to the management group that contains the existing cluster.

Storage nodes and cluster capacity

Be certain that the capacity of the storage node you add to the cluster matches, or is close to, the capacity of those already in the cluster. All storage nodes in a cluster operate at a capacity equal to that of the smallest capacity storage node. If you add a storage node with a smaller capacity, the capacity of the entire cluster will be reduced.

While you can mix storage nodes with different RAID levels in a cluster, note that the capacity limitation applies to the available capacity as determined by RAID, not the raw disk capacity.

Example

If you have three storage nodes, two of which have a capacity of 1 terabyte, and one of which has a capacity of 2 TB, all three storage nodes operate at the 1 TB capacity.

Adding storage to a cluster

1. Select the cluster in the navigation window.
2. Click Cluster Tasks and select Edit Cluster.

If there are no storage nodes in the management group available to add to the cluster, the Add Nodes button will be greyed out.

3. Click Add Nodes.
4. Select one or more storage nodes from the list.
5. Click OK.
6. Click OK again in the Edit Clusters window.

A confirmation message opens, describing the restripe that happens when a storage node is added to a cluster.

7. Click OK to finish adding the storage node to the cluster.

Removing a storage node from a cluster

You can remove a storage node from a cluster only if the cluster contains sufficient storage nodes to maintain the existing volumes and their replication level. See “[Guide for volumes](#)” on page 238 for details about editing volumes.

1. In the Edit Cluster window, select a storage node from the list.
2. Click Remove Nodes.

In the navigation window, that storage node moves out of the cluster, but remains in the management group.

3. Click OK when you are finished.



NOTE:

Changing the order of the storage node list causes a full cluster restripe.

Changing or removing the virtual IP

Anytime you add, change or remove the virtual IP address for iSCSI volumes, you are changing the configuration that servers are using. You should re-balance the iSCSI sessions after making the change.

Preparing servers

- Quiesce any applications that are accessing volumes in the cluster.
- Log off the active sessions in the iSCSI initiator for those volumes.

Changing the virtual IP address

1. In the Edit Cluster window, click the iSCSI tab to bring it to the front.
2. Select the VIP you want to change.
3. Change the information in the Edit VIP and Subnet Mask window.
4. Click OK to return to the Edit Cluster window.

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.
A confirmation message opens.
3. Click OK to confirm the deletion.

Finishing up

1. Click OK when you are finished changing or removing the VIP.
2. Reconfigure the iSCSI initiator with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

Changing or removing an iSNS server

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.

Changing an iSNS server

1. Select the iSNS server to change.
2. Click Edit.
The Edit iSNS Server window opens.
3. Change the IP address.
4. Click OK.

Deleting an iSNS server

1. Select the iSNS server to delete.

2. Click Delete.
A confirmation message opens.
3. Click OK.

Finishing up

1. Click OK when you are finished changing or removing an iSNS server.
2. Reconfigure the iSCSI initiator with the changes.
3. Reconnect to the volumes.
4. Restart the applications that use the volumes.

Troubleshooting a cluster

Auto Performance Protection monitors individual storage node health related to performance issues that affect the volumes in the cluster.

Repairing a storage node provides a way to replace a failed disk in a storage node and minimize the time required to bring the storage node back to normal operation in the cluster with fully resynchronized data.

Auto Performance Protection

If you notice performance issues in a cluster, a particular storage node may be experiencing slow I/O performance, overload or latency issues. You can identify whether Auto Performance Protection is operating by checking the storage server status on the storage node Details tab.

Auto Performance Protection is indicated by two unique statuses reported on the Details tab. You will also receive alert notifications on these statuses.

- **Storage Server Overloaded.** The Overloaded status indicates that operations to the storage node are completing too slowly. During the overloaded state, volume availability is maintained while the storage node is quarantined in the cluster. While the storage node 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 node is returned to active duty and resynced with the data that has changed since its quarantine. Volumes that depend on this storage node will then show “Resyncing” on the volume Details tab.
- **Storage Server Inoperable.** The Inoperable status indicates that the storage node is unable to repair the slow I/Os, which may indicate a potential hardware problem. Volumes that depend on this storage node are unavailable. For information about how to determine volume availability, see the section “[Determining volume and snapshot availability](#)” on page 51.
Rebooting the storage node 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.

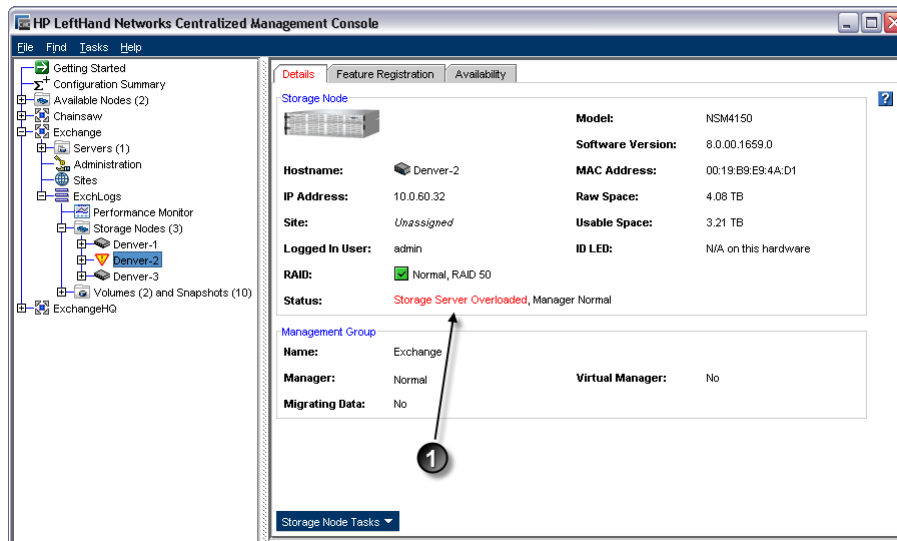
Auto Performance Protection and other clusters

Auto Performance Protection operating on a storage node in one cluster will not affect performance for other clusters in the management group.

Checking storage node status

You can easily identify whether Auto Performance Protection is active on a storage node in a cluster with performance issues.

1. Select the affected storage node in the navigation window.
The storage node icon will be blinking in the navigation tree.
2. Check the Status line on the Details tab.



1. Status line

Figure 99 Checking the storage node status on the Details tab

If status is Storage Server Overloaded

Wait up to 10 minutes and check the status again. The status may return to Normal and the storage node will be resyncing.

If status is Storage Server Inoperable

Reboot the storage node 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 node

Repairing a storage node allows you to replace a failed disk in a storage node that contains volumes configured for 2-way or 3-way replication, and trigger only one resync of the data, rather than a complete restriping. Resyncing the data is a shorter operation than a restripe.

Prerequisites

- Volume must have 2-way or 3-way replication.
- Storage node must have the blinking red and yellow triangle in the navigation window.
- If the storage node is running a manager, stopping that manager must not break quorum.

How repair storage node works

Using Repair Storage Node to replace a failed disk includes the following steps:

- Using Repair Storage Node from the Storage Node Tasks menu to remove the storage node from the cluster
- Replacing the disk in the storage node
- Returning the storage node to the cluster


Because of the replication level, removing and returning the storage node to the cluster would normally cause the remaining storage nodes in the cluster to restripe the data twice—once when the storage node is removed from the cluster and once when it is returned.

The Repair Storage Node command creates a placeholder in the cluster, in the form of a “ghost” storage node. This ghost storage node keeps the cluster intact while you remove the storage node, replace the disk, configure RAID, and return the storage node to the cluster. The returned storage node only has to resynchronize with the other 2 storage nodes in the cluster.

Using the repair storage node command

When a storage node in a cluster has a disk failure, the navigation window displays the storage node and the cluster with a blinking triangle next to them in the tree. An alert appears in the alert window, and the Status label in the tab window shows the failure.

1. If the storage node is running a manager, stop the manager. See “[Stopping managers](#)” on page 181.
2. Right-click the storage node and select Repair Storage Node.

3. From the Repair Storage Node window, select the item that describes the problem you want to solve. Click More for more detail about each selection.
 - Repair a disk problem
If the storage node has a bad disk, be sure to read [“Replacing a disk”](#) on page 80 before you begin the process.
 - Storage Node problem
Select this choice if you have verified that the storage node must be removed from the management group to fix the problem. For more information about using Repair Storage Node with a disk replacement, see [“Replacing disks”](#) on page 328.
 - Not sure
This choice allows you to confirm whether the storage node has a disk problem by taking you directly to the Disk Setup window so that you can verify disk status. As in the first choice, be sure plan carefully for a disk replacement.
4. Click OK.
The storage node leaves the management group and moves to the Available Nodes pool. A placeholder, or “ghost” storage node 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 node 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”](#) on page 80.
6. Return the repaired storage node to the management group.
The ghost storage node remains in the cluster.



NOTE:

The repaired storage node will be returned to the cluster in the same place it originally occupied to ensure that the cluster resyncs, rather than restripes. See [Chapter 24](#) on page 347 for definitions of restripe and resync.

7. [Optional] Start a the manager on the repaired storage node.

To return the repaired storage node to the cluster

1. Right-click the cluster and select Edit Cluster window.

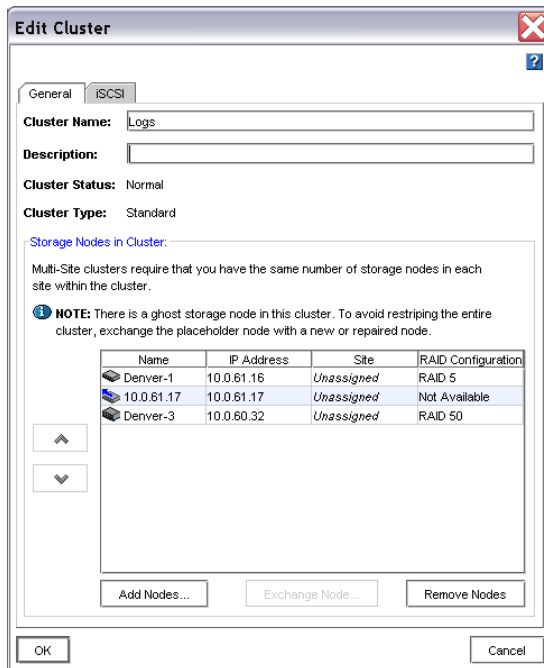


Figure 100 Exchanging ghost storage node

2. Select the ghost storage node (the IP address in the list) and click Exchange Node.

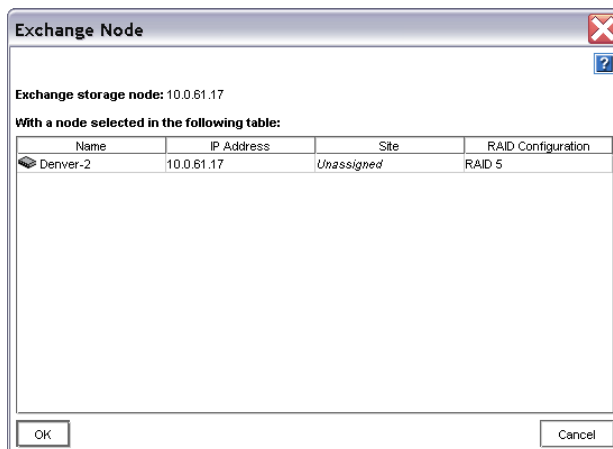


Figure 101 Replacing the repaired storage node

3. Select the repaired storage node to exchange for the ghost storage node and click OK.
The storage node returns to its original position in the cluster and volumes in the cluster proceed to resync.

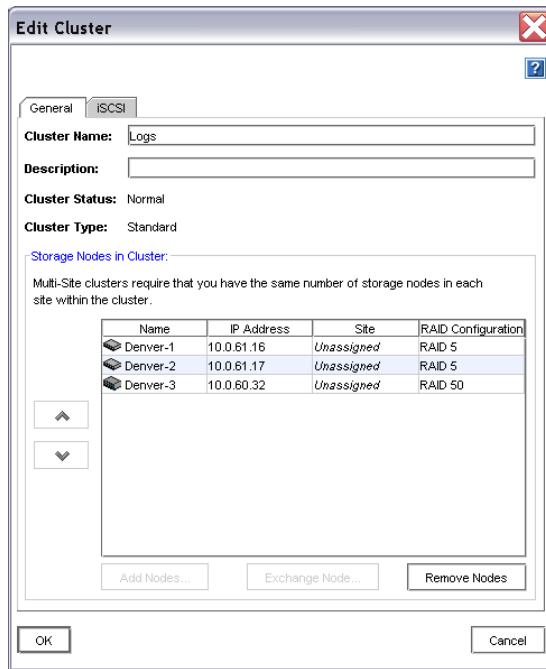


Figure 102 Repaired storage node returns to proper place in cluster

Deleting a cluster

Volumes and snapshots must be deleted or moved to a different cluster before you can delete the cluster. For more information, see [“Deleting a volume”](#) on page 243 and [“Deleting a snapshot”](#) on page 261.

12 Provisioning storage

The SAN/iQ software uses 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 and related policies, 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 platforms and the features of the SAN/iQ software.

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 nodes and the RAID level on them.

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 nodes which is determined by the disk capacity and the RAID level configured on the storage nodes.
See [“Planning the RAID configuration”](#) on page 67.
- The second factor is the effect of the replication level of the volumes and snapshots.
See [“Planning data replication”](#) on page 223.
- The third factor is the snapshot configuration, including schedules and retention policies.
See [“Managing capacity using volume size and snapshots”](#) on page 226.
- 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 cluster, allows you to free up space on the application cluster more rapidly.
See the chapter [“Understanding and Planning Remote Copy”](#) in the *Remote Copy User Manual*.

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.

Best practices

To take full advantage of the features of the HP LeftHand Storage Solution, use the appropriate combination of RAID, volume and snapshot replication, and schedules to snapshot a volume and related retention policies.

Table 44 Recommended SAN configurations for provisioning storage

RAID	Replication Level
RAID0	2-Way or 3-Way
RAID10	2-Way
RAID5	2-Way
RAID6	2-Way

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 45 Volume provisioning methods

Method	Settings
Full provisioning	Volume size x replication level = amount of space allocated on the SAN
Thin provisioning	Volume size > 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. Use thin provisioning when the application that is writing to the volume is effective at reusing disk space. However, the SAN/iQ software warns you that the cluster is nearly full. You always know that a thin volume may risk a write failure.

The SAN/iQ software allocates space as needed. However, thin provisioning carries the risk that an application server will fail a write because the SAN has run out of disk space.

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](#)” on page 234.

Planning data replication

Data replication creates redundant copies of a volume on the SAN. You can create up to four copies using 4-Way replication. Because these copies reside on different storage nodes, replication levels are tied to the number of available storage nodes in a cluster.

The SAN/iQ software and the HP LeftHand Centralized Management Console provide flexibility through two features when you are planning data replication.

- Replication levels allow you to choose how many copies of data you want to keep on the cluster.
- Replication priority allows you to choose whether data availability or data redundancy is more important in your configuration.

Replication level

Four replication levels are available depending upon the number of available storage nodes in the cluster. The level of replication you choose also affects the Replication Priority you can set.

Table 46 Setting a replication level for a volume

With this number of available storage nodes in cluster	Select this replication level	For this number of copies
1	<ul style="list-style-type: none"> • None 	<ul style="list-style-type: none"> • 1 copy of data in the cluster. No replica is created.
2	<ul style="list-style-type: none"> • None • 2-Way 	<ul style="list-style-type: none"> • 1 copy of data in the cluster. No replica is created. • 2 copies of data in the cluster. One replica created.
3	<ul style="list-style-type: none"> • None • 2-Way • 3-Way 	<ul style="list-style-type: none"> • 1 copy of data in the cluster. No replica is created. • 2 copies of data in the cluster. One replica created • 3 copies of data in the cluster. Two replicas created.
More than 3	<ul style="list-style-type: none"> • None • 2-Way • 3-Way • 4-Way 	<ul style="list-style-type: none"> • 1 copy of data in the cluster. No replica is created. • 2 copies of data in the cluster. One replica created • 3 copies of data in the cluster. Two replicas created. • 4 copies of data in the cluster. 3 replicas created.

How replication levels work

When you choose 2-Way, 3-Way, or 4-Way replication, data is written to either 2, 3, or 4 adjacent storage nodes in the cluster. The system calculates the actual amount of storage resources needed if the replication level is greater than None.

2-way replication

A cluster with 4 storage nodes is configured for 2-Way replication. There have been 4 writes to the cluster. [Figure 103](#) illustrates the write patterns on the 4 storage nodes.




Figure 103 Write patterns in 2-Way replication

3-Way replication

A cluster with 4 storage nodes is configured for 3-Way replication. There have been 4 writes to the cluster. [Figure 104](#) illustrates the write patterns on the 4 storage nodes.

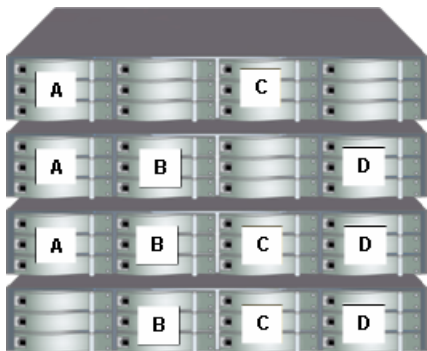


Figure 104 Write patterns in 3-Way replication

4-Way replication

A cluster with 4 storage nodes is configured for 4-Way replication. There have been 4 writes to the cluster. [Figure 105](#) illustrates the write patterns on the 6 storage nodes.




Figure 105 Write patterns in 4-Way replication

Replication priority

Set the replication priority according to your requirements for data availability or data redundancy for the volume.

Redundancy Mode

Choose the redundancy mode if you require that the volume be replicated in order to be available. The redundancy mode ensures fault-tolerance.

Availability Mode

Choose the availability mode (which is the default) if you want your data to be available even if it cannot be replicated. The availability mode ensures that data may remain available to servers even if a storage node becomes unavailable.

Table 47 Storage node availability and volume access by replication level and priority setting

Volume is available to a server with				
a priority setting of:	and a replication level of:			
	None	2-Way	3-Way	4-Way
Availability	All storage nodes must be up.	1 of every 2 adjacent storage nodes must be up. Adjacent storage nodes are those that are adjacent in the cluster.	1 of every 3 adjacent storage nodes must be up.	1 of every 4 adjacent storage nodes must be up.
Redundancy	N/A	All storage nodes must be up.	2 of every 3 adjacent storage nodes must be up.	2 of every 4 adjacent storage nodes must be up.

△ CAUTION:

A management group with 2 storage nodes and a Failover Manager is the minimum configuration for automated fault tolerant operation. Although the SAN/iQ software allows you to configure 2-way replication on 2 storage nodes, this does not guarantee data availability in the event that one storage node becomes unavailable, due to the communication requirements between managers. See “[Managers overview](#)” on page 171.

Best practice for setting replication levels and redundancy modes

For mission-critical data and using a 3-node cluster, choose 3-Way or 4-Way replication and redundancy priority. This configuration sustains the first fault and ensures that the volume is redundant and available.

If your volumes contain critical data, configure them for 2-Way replication and a priority of redundancy.

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. Prior to this release, snapshots were full provisioned.

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 [Chapter 14](#) on page 245.

Managing capacity using volume size and snapshots

How snapshots are created

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. The effect of this data migration can be seen in the Max Provisioned Space and Used Space columns on the Volume Usage tab of the cluster. See "[Ongoing capacity management](#)" on page 227 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 node 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](#)" on page 174 .

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 nodes in the cluster. This information is presented in a series of tab windows presented at the cluster level.

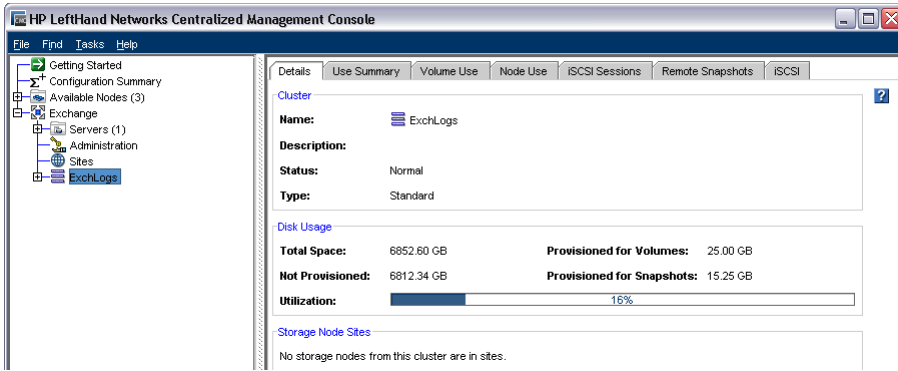


Figure 106 Cluster tab view

Cluster use summary

The Use Summary window presents information about the total space available in the cluster, the amount of space provisioned for volumes and snapshots, and how much of that space is currently used by volumes and snapshots.

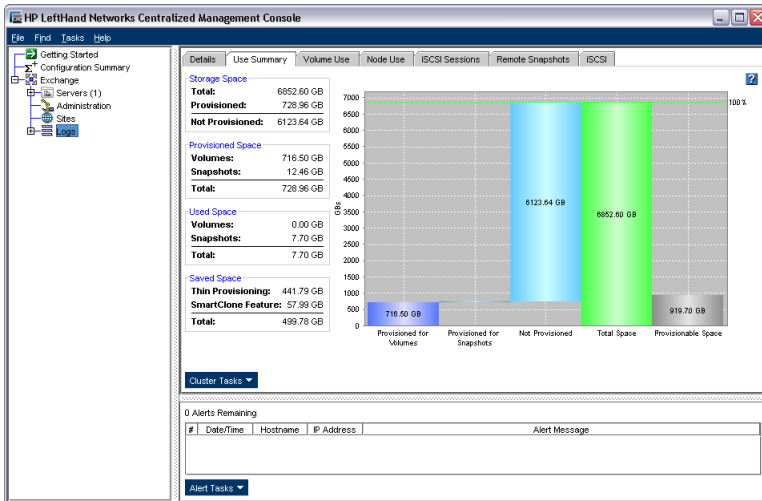


Figure 107 Reviewing the Use Summary tab

In the Use Summary window, the Storage Space section reflects the space available on the storage nodes in the cluster. Storage space is broken down as shown in Table 48.

Table 48 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.

Category	Description
Not provisioned	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.
Provisioned space	
Volumes	Amount of space allocated for volumes. For fully provisioned volumes, this is the size x the replication level. For thinly provisioned volumes, the amount of space allocated is determined by the system.
Snapshots	Amount of space allocated for snapshots and temporary space, if required. This value is zero until at least one snapshot has been created. If all snapshots are deleted, this value returns to zero.
Total	Sum of the space allocated for volumes, snapshots, and temporary space.
Used space 1	
Volumes	Actual amount of space used by volumes.
Snapshots	Actual amount of space used by snapshots, including temporary space.
Total	Total of space used by volumes and snapshots. For more information, see Measuring disk capacity and volume size (page 233).
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. 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 for volumes	Amount of space allocated for volumes. For fully provisioned volumes, this is the size x the replication level. For thinly provisioned volumes, the amount of space allocated is determined by the system.
Provisioned for snapshots	Amount of space allocated for snapshots and temporary space, if required. This value is zero until at least one snapshot has been created. If all snapshots are deleted, this value returns to zero.
Not provisioned	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 space	Combined space available in the cluster for storage volumes and snapshots.

Category	Description
Max provisioned space	Total space that volumes and snapshots can grow to fill. Note: In the case of overprovisioning, this value can exceed the physical capacity of the SAN.

¹The used space decreases when you delete volumes, snapshots or temporary space from the SAN. The Cluster Summary Used Space can also decrease when a volume is moved. Deleting files or data from client applications does not decrease the used space.

Volume use summary

The Volume Use window presents detailed information about the volume characteristics that affect the utilization of the cluster. The table lists the volumes and snapshots, and the space and utilization totals for the cluster.

Table 49 Information on the Volume Use tab

Category	Description
Name	Name of the volume, snapshot or cluster.
Size	Size of the volume or snapshot presented to the server. In the case of snapshots, the size is automatically determined, and is set to the size of the parent volume at the time the snapshot was created.
Replication level	Choices include None, 2-Way, 3-Way, or 4-Way. Snapshots inherit the replication level of the parent volume.
Provisioning type	<p>Volumes can be either full or thin provisioned. Snapshots are always thin provisioned, unless you are viewing a fully provisioned snapshot in SAN/iQ software version 6.6 or earlier. The Provisioning Type column also details space saving options for the different types of volumes and snapshots you can create on the SAN, as shown in Figure 108 on page 231. The space calculations take into account both the type of volume and the replication level of the volume or snapshot. 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 to see the space saved number decrease as data on the volume increases. Full provisioning allocates the full amount of space for the size of the volume. Reclaimable space is the amount of space that would be on the SAN if this fully provisioned volume was changed to thinly provisioned. <p>The totals at the cluster level, shown at the bottom of the list, show the total for both saved and reclaimable space.</p>

Category	Description
----------	-------------

Provisioned space

The provisioned space is the amount of space reserved for data on the SAN. Temporary space is space used by applications and operating systems that need to write to a snapshot when they access it. [Figure 109](#) on page 232 shows temp space that can be deleted or converted to a volume.

- Full provisioned volumes display the entire amount of allocated space in this column, which is the volume size times the replication level. For example, 10 GB size x 2-Way replication results in 20 GB of provisioned space.
- Thin provisioned volumes allocate a fraction of the total amount of space planned. The provisioned space increases as needed, up to the maximum provisioned space or until the cluster is full.
- Snapshots are automatically thin provisioned. The provisioned space is that which is allocated when the snapshot is created. The amount of provisioned space can change as snapshots are deleted.
- The temporary space is equal to the size of the snapshot. For example, if a snapshot size equals 2 GB, the temporary space is also 2 GB.

Max provisioned space

The total amount of space which can be allocated for the volume, assuming there is enough space in the cluster.

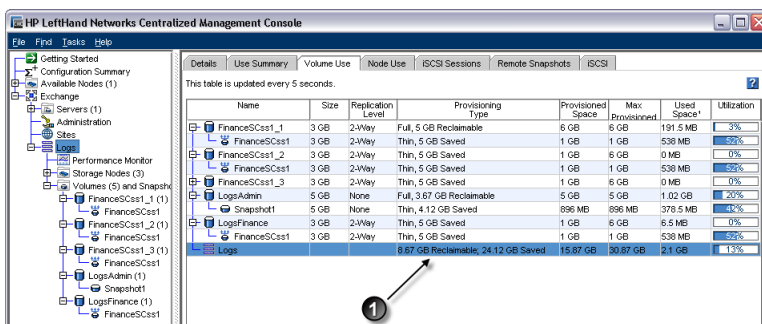
- Full provisioned volumes - this is the same as provisioned space described above.
- Thin provisioned volumes - this total reflects the size of the volume times the replication level.
- Snapshots - this value is the same as snapshot provisioned space, unless there is temporary space for the snapshot. In that case, the temporary space is also reflected in this total.

Used space

Amount of space used by actual data in the volume or snapshot. Used space only decreases when you delete volumes, snapshots or temporary space from the SAN. The total of cluster used space may decrease if volumes are deleted or moved to a different cluster. Deleting files or data from client applications does not decrease the used space. For more information, see [“Measuring disk capacity and volume size”](#) on page 233.

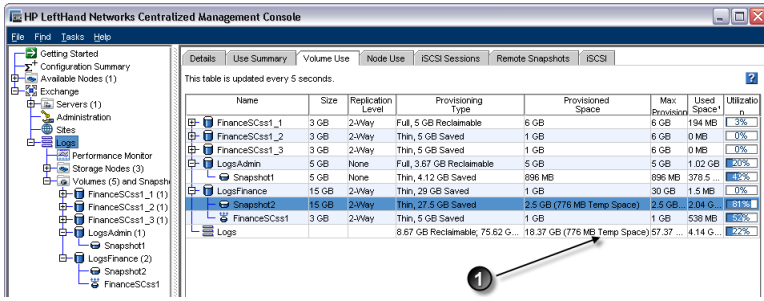
Utilization

Percentage of the Max Provisioned Space that has been written to. This value is calculated by dividing the Used Space by the Max Provisioned Space.



1. Space saved or reclaimable displayed here

Figure 108 Viewing the space saved or reclaimable in the Volume Use tab



1. Temp space can be deleted or converted to a volume

Figure 109 Provisioned space shows temp space used

Node use summary

The Node Use window presents a representation of the space provisioned on the storage nodes in the cluster.

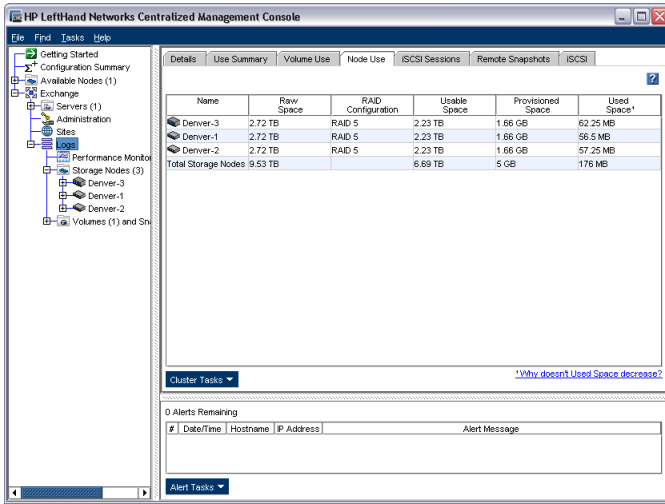



Figure 110 Viewing the Node Use tab

Table 50 Information on the Node Use tab

Category	Description
Name	Hostname of the storage node.
Raw space	Total amount of disk capacity on the storage node. The raw space column also shows the effect of putting storage nodes of different capacities in the same cluster. For example, in Figure 111 on page 233, Denver-3 shows the raw space value in bold and the note that some storage space is stranded. Stranded storage occurs when storage nodes in a cluster are not all the same capacity. Storage nodes with greater capacity will only operate to the capacity of the lowest capacity storage node in the cluster. The remaining capacity is considered stranded and the Raw Space column shows a bolded value for the higher-capacity storage nodes. The stranded storage space can be reclaimed by equalizing the capacity of all the nodes in the cluster.
RAID configuration	RAID level configured on the storage node.

Category	Description
Usable space	Space available for storage after RAID has been configured.
Provisioned space	Amount of space allocated for volumes and snapshots.
Used space	Amount of space consumed by volume or snapshot data on this storage node. This value never decreases, though the actual space available may grow and shrink as data is manipulated through a file system, if one is configured on the volume. For more information, see “Measuring disk capacity and volume size” on page 233.



1. Denver-3 with stranded storage space

Figure 111 Stranded storage in the cluster

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 51 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 re-orders 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 on-line and available to users.

Increasing the volume size in other environments

Some environments use alternative tools, such as Dell Array Manager and VERITAS Volume Manager. Both of these disk management tools use a utility called `Extpart.exe` instead of `Diskpart.exe`. `Extpart.exe` commands are similar to those of `Diskpart.exe`. The only major difference is that instead of selecting the volume number, as in `Diskpart.exe`, you select the drive letter instead. `Extpart.exe` and corresponding documentation can be downloaded from www.dell.com.

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
- Deleting snapshot temporary space

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](#)” on page 233. For file-level capacity management, use application or file system-level tools.

Snapshot temporary space

When you mount a snapshot, additional space can be created in the cluster for use by applications and operating systems that need to write to the snapshot when they access it. This additional space is called the temporary space. For example, MS Windows performs a write when the snapshot is mounted via iSCSI. Microsoft Volume Shadow Copy Service (VSS) and other backup programs write to the snapshot when backing it up.

The amount of temporary space initially provisioned on the SAN is minimal. However, if you do write data to the snapshot, it goes to the temporary space, which then grows as necessary to accommodate the amount of data written. You can see how much temporary space is being used for a snapshot on the Volume Use tab in the Cluster tab window.

Managing snapshot temporary space

You can manage the temporary space two ways — delete it or convert it to a volume.

Delete the space to free up space on the cluster

The additional temporary space is deleted when the snapshot is deleted. If you need to free up the extra space before the snapshot is deleted, you can do so manually in the CMC or through your snapshot scripts. The next time an application or operating system accesses the snapshot, a new, empty temporary space is created.

For instructions to delete snapshot temporary space, see [“Delete the temporary space”](#) on page 253.

Convert temporary space to a volume

If you have written data to a mounted snapshot and you need to permanently save or access that data, you can convert the temporary space to a volume. That volume will contain the original snapshot data plus any additional data written after the snapshot was mounted.

For instructions to convert snapshot temporary space, see [“Convert the temporary space”](#) on page 253.

13 Using volumes

A volume is a logical entity that is made up of storage on one or more storage nodes. 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 nodes.

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 [Chapter 17](#) on page 289.

Prerequisites

Before you create a volume, you must have created a management group and at least one cluster. For more information, see the following:

- [Chapter 9](#) on page 171
- [“Creating additional clusters”](#) on page 209

Planning volumes

Planning volumes takes into account multiple factors.

- How many volumes do you need?
- What type of volume are you creating - primary or remote?
- What size do you want the volume to be?
- Do you plan to use snapshots?
- Do you plan to use data replication?
- Do you plan to grow the volume or to keep 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”](#) on page 174.

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 *Remote Copy User Manual* 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 [Chapter 15](#) on page 263.

Guide for volumes

When creating a volume, you define the following characteristics.

Table 52 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” on page 34 for more information.
Description	Both	[Optional] A description of the volume.
Size	Primary	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 nodes 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 Chapter 14 on page 245. Remote volumes contain no data and so do not have a size.
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.
Replication Level	Both	Default value = 2. The number of copies of the data to create on storage nodes in the cluster. The replication level is, at most, the number of storage nodes in the cluster or 4, whichever is smaller. If you select a replication level of None, the replication priority is not used. See “Planning data replication” on page 223.

Volume characteristic	Configurable for Primary or Remote Volume	What it means
Replication Priority	Both	<p>Default value = Availability</p> <ul style="list-style-type: none"> • Availability — These volumes remain available as long as at least one storage node out of every n (n = replication level-adjacent storage nodes) remains active. When the unavailable storage node returns to active status in the cluster, then the volume re-synchronizes across the replicas. • Redundancy — This setting ensures that the volume becomes unavailable if it cannot maintain 2 replicas. For example, if 2-way replication is selected, and a storage node in the cluster becomes unavailable, thereby preventing 2-way replication, the volume becomes unavailable until the storage node is again available.
Type	Both	<p>Default value = Primary</p> <ul style="list-style-type: none"> • Primary volumes are used for data storage. • Remote volumes are used for configuring Remote Copy for business continuity, backup and recovery, or data mining/migration.
Provisioning	Primary	<p>Default value = Full</p> <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 SAN/iQ software allocates space as needed. However, thin provisioning carries the risk that 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 nodes 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.

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)
- Replication level
- Replication priority
- Volume type
- Provisioning

Descriptions of these characteristics are found in [Table 52](#) on page 238.

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.



NOTE:

The system automatically factors replication levels into the settings. For example, if you create a 500 GB volume and the replication level is 2, the system automatically allocates 1,000 GB for the volume.

Editing a volume

When editing a primary volume, you can change the description, size, and advanced characteristics such as the cluster, replication level, replication priority, 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 53 Requirements for changing volume characteristics

Item	Requirements for Changing
Description	Must be from 0 to 127 characters.

Item	Requirements for Changing
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 nodes and unallocated space for the size and replication 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>
Replication Level	The cluster must have sufficient storage nodes and unallocated space to support the new replication level. For example, you just added more storage to a cluster and have more capacity. You decide to change the replication level for a volume from 0 to 2 to ensure you have redundancy for your data.
Replication Priority	To change the replication priority, the replication level must support the change. You can always go from Redundancy to Availability. However, you cannot go from Availability to Redundancy unless a sufficient number of storage nodes in the cluster are available. For example, if you have 2-way replication with 3 storage nodes in the cluster, you can change from Availability to Redundancy if all the storage nodes in the cluster are available. You can use Redundancy to ensure data integrity if you know you are going to take a cluster offline. Redundancy ensures that when any one storage node goes offline then the volume becomes unavailable in order to protect the data.
Size	<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 node to the cluster <p>To decrease the size of the volume</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.
The Edit Volume window opens.

Changing the volume description

1. In the Description field, edit the description.

2. Click OK when you are finished.

Changing the 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 LeftHand DSM for MPIO, log off the volumes from the server, add the VIP or the individual IP addresses of the storage nodes 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 replication level

1. In the Replication Level drop-down list, select the level of replication you want.
2. Click OK when you are finished.

Changing the replication priority

1. Select the replication priority you want.
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.

△ 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.

Making an unavailable redundancy volume available

If a storage node becomes unavailable and needs to be repaired or replaced, and a replicated volume that is configured for redundancy becomes unavailable to servers, the following procedure allows you to return the volume to fully operational status.

1. Stop any applications that are accessing the volume and log off all associated iSCSI sessions.
2. Select the volume in the navigation window.
3. Right-click and select Edit Volume.
4. On the Advanced tab, change the Replication Priority from Redundancy to Availability.
The Replication Level must be 2 or greater.

5. Use the Repair Storage Node procedure, described in “[Repairing a storage node](#)” on page 216.
6. Reconnect the iSCSI sessions and restart the applications that access the volume.

Deleting a volume

Delete a volume to remove that volume’s data from the storage node 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](#)” on page 272 and “[Shared snapshot](#)” on page 274.

△ **CAUTION:**

Deleting a volume permanently removes that volume’s data from the storage node.

Prerequisites

- Stop any applications that are accessing the volume and log off all associated iSCSI sessions

New in release 8.0

Deleting a volume automatically deletes all associated snapshots except those that are clone points or shared snapshots as part of a SmartClone volume configuration. Releases before 8.0 required all associated snapshots to be deleted manually before deleting the volume.

To delete the volume

1. In the navigation window, select the volume you want to delete.
The Volume tab window opens.
2. Click Volume Tasks and select Delete Volume.
A confirmation window opens.
3. Click OK.
The volume is removed from the cluster.

To delete multiple volumes

1. In the navigation window, select Volumes and Snapshots.

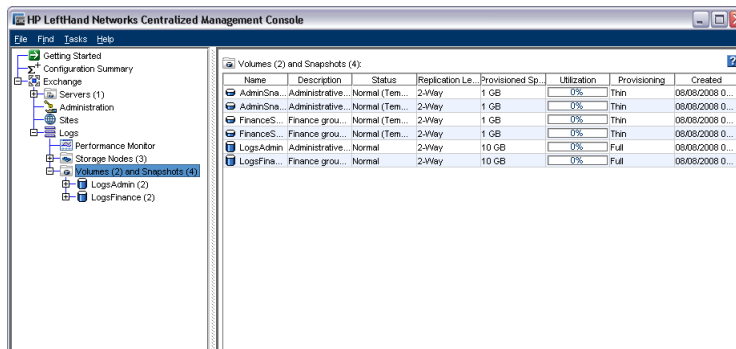


Figure 112 Viewing multiple volumes and snapshots

2. Shift+click or Ctrl+click to select the volumes and snapshots you want to delete.

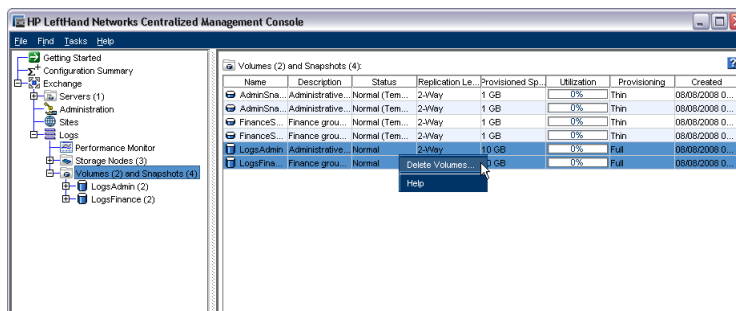


Figure 113 Deleting multiple volumes in one operation

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.

14 Using snapshots

Using snapshots overview

Snapshots are a copy of a volume for use with backup and other applications. Snapshots are one of the following types:

- 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 LeftHand P4000 VSS Provider (VSS Provider). For more information, see [“Requirements for application-managed snapshots”](#) on page 248.
- Point-in-time consistent — 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 the data.

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.

Prerequisites

Before you create a snapshot, you must create a management group, a cluster, and a volume to receive it. Use the Management Groups, Clusters and Volumes wizard to create them.

For information, see

- [“Creating a management group”](#) on page 177
- [“Creating additional clusters”](#) on page 209
- [Creating a volume](#) (page 239)
- [“The effect of snapshots on cluster space”](#) on page 226

Using snapshots

You create snapshots from a volume on the cluster. At any time you can roll a volume back to a specific snapshot. You can mount a snapshot to a different server and recover data from the snapshot to that server. You can also create a SmartClone volume from a snapshot.

Snapshots can be used for these cases:

- Source for creating backups
- Data or file system preservation before upgrading software
- Protection against data deletion
- 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 [Chapter 15](#) on page 263 .

Single snapshots versus scheduled snapshots

Some snapshot scenarios call for creating a single snapshot and then deleting it when it is no longer needed. Other scenarios call for creating a series of snapshots up to a specified number or for a specified time period, after which the earliest snapshot is deleted when the new one is created (snapshots created from a schedule).

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.

Guide for snapshots

Review “[Planning volumes](#)” on page 237 to ensure that you configure snapshots correctly. When creating a snapshot, you define the following characteristics or options.

Table 54 Snapshot characteristics

Snapshot Parameter	What it means
Application-managed Snapshot	This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot. This option requires the use of the VSS Provider. For more information, see “ Requirements for application-managed snapshots ” on page 248. If the VSS Provider is not installed, SAN/iQ will let you create a point-in-time consistent snapshot (not using VSS).
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 ” on page 34 for information about this naming convention. The following are illegal characters: , ' " ; : =.
Description	(Optional) A description of the snapshot.
Assign and Unassign Servers	(Optional) Configure server access to the snapshot.

Planning snapshots

When planning to use snapshots, consider their purpose and size. If you are planning to use schedules to snapshot a volume, see “[Storage nodes and cluster capacity](#)” on page 211 and [Table 55](#) on page 246 for approximate data change rates for some common applications.

Table 55 Common applications’ daily change rates

Application	Daily Change Rates
Fileshare	1 - 3%
Email/Exchange	10 - 20%

Application	Daily Change Rates
Database	10%



NOTE:

When considering the size of snapshots in the cluster, remember that the replication level of the volume is duplicated in the snapshot.

Source volumes for tape backups

Best practice

Plan to use a single snapshot and delete it when you are finished. Consider the following question in your planning.

- Is space available on the cluster to create the snapshot?

Data preservation before upgrading software

Best practice

Plan to use a single snapshot and delete it when you are finished. Consider the following questions in your planning.

- Is space available on the cluster to create the snapshot?

Automated backups

Best practice

Plan to use a series of snapshots, deleting the oldest on a scheduled basis. Consider the following questions in your planning.

- 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](#)” on page 246 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.

Planning how many snapshots

For information about the recommended maximum number of volumes and snapshots that can be created in a management group, see “[Configuration summary overview](#)” on page 174 and [Chapter 9](#) on page 171.

Creating a snapshot

Create a snapshot to preserve a version of a volume at a specific point in time. For information about snapshot characteristics, see [“Guide for snapshots”](#) on page 246.

1. Log in to the management group that contains the volume for which you want to create a new snapshot.
2. Right-click on the volume and select New Snapshot.
3. If you want to use VSS to quiesce the application before creating the snapshot, select the Application-Managed Snapshot option.

This option requires the use of the VSS Provider. For more information, see [“Requirements for application-managed snapshots”](#) on page 248. If the VSS Provider is not installed, SAN/iQ will let you create a point-in-time consistent snapshot (not using VSS).

This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot.

The system fills in the Description and Servers fields automatically.

4. Type a name for the snapshot.
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.

Requirements for application-managed snapshots

For single snapshots, you can create application-managed snapshots. Application-managed snapshots use the VSS Provider to quiesce VSS—aware applications before creating the snapshot. The following are required for application-managed snapshots:

- SAN/iQ version 8.0 or later
- CMC or CLI latest update
- HP LeftHand P4000 Solution Pack, specifically the HP LeftHand P4000 VSS Provider (latest update) installed on the application server (refer to the HP LeftHand P4000 Windows Solution Pack User Manual)
- Management group authentication set up for the VSS Provider (refer to the HP LeftHand P4000 Windows Solution Pack User Manual)
- Application on the server that is VSS-aware
- Server set up in SAN/iQ with iSCSI connection (see [Chapter 17](#) on page 289)

Creating an application-managed snapshot using SAN/iQ is the same as creating any other snapshot. However, you must select the Application-Managed Snapshot option. For information about creating snapshots, see [“Creating a snapshot”](#) on page 248.

Creating snapshots for volume sets

The snapshot creation process for application-managed snapshots differs only when an application has associated volumes. Associated volumes are two or more volumes used by an application (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.

When you create an application-managed snapshot of a volume in a volume set, the CMC recognizes that the volume is part of a volume set. SAN/iQ then prompts you to create a snapshot for each volume in the volume set. This creates a snapshot set that corresponds to the volume set. A future release will identify snapshot sets in the CMC.

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 procedure below assumes that you select a volume that is part of a volume set for the snapshot.

1. Log in to the management group that contains the volume for which you want to create a new snapshot.
2. Right-click on the volume and select New Snapshot.
3. Select the Application-Managed Snapshot option.

This option requires the use of the VSS Provider. For more information, see [“Requirements for application-managed snapshots”](#) on page 248.

This option quiesces VSS-aware applications on the server before SAN/iQ creates the snapshot.

The system fills in the Description and Servers fields automatically. You cannot edit them.

4. Type a name for the snapshot.
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 for each snapshot.

NOTE:

Be sure to leave the Application-Managed Snapshots option selected. This option maintains the association of the volumes and snapshots and quiesces the application before creating the snapshots. If you deselect the option, the system creates a point-in-time consistent snapshot of each volume listed.

-
7. Click Create Snapshots to create a snapshot of each volume.

The snapshots all display in the CMC. A future release will identify snapshot sets in the CMC.

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.

The snapshot Details tab refreshes.

Mounting or accessing 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”](#) on page 259.
- 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”](#) on page 253 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. For information about creating and using servers, see [Chapter 17](#) on page 289.

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 an application-managed snapshot as a volume, use diskpart.exe to change the resulting volume's attributes.

For more information, see [“Making an application-managed snapshot available”](#) on page 250.

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 an application-managed snapshot available

If you do any of the following using an 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 an application-managed snapshot available on a stand-alone server

Use this procedure to make an 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 `detaildisk`.
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 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 LeftHand Storage Solution 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 an 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 `detaildisk`.

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, refer to 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 LeftHand Storage Solution 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

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

1. Right-click the snapshot for which you want to save the additional data.
2. Select Convert Temporary Space from the menu.
3. Type 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”](#) on page 257.

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 an application-managed snapshot available”](#) on page 250.

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.

Prerequisite

- Stop any applications that are accessing the snapshot and log off all related iSCSI sessions.

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.

1. In the navigation window, select the 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.

Creating a schedule to snapshot a volume

You can schedule recurring snapshots of a volume. Recurring snapshots of volumes can be scheduled in a variety of frequencies and with a variety of retention policies. You can schedule a snapshot of a volume every 30 minutes or more, and retain up to 50 snapshots.

If you need to, you can pause and resume any schedule to snapshot a volume. Currently, you cannot create scheduled, application-managed snapshots from SAN/iQ. This function will be available in a future release.

 **NOTE:**

Scripting 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

- Schedules to snapshot a volume require particular attention to capacity management. See “[Understanding how the capacity of the SAN is used](#)” on page 221.
- 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 nodes 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 56 Requirements for scheduling 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” on page 226. 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” on page 174.</p>

Creating schedules to snapshot a volume

You can create one or more schedules to snapshot a volume. For example, your backup and recovery plan might include three schedules: one schedule for daily snapshots, retained for seven days; the second schedule for weekly snapshots, retained for four weeks; the third schedule for monthly snapshots, retained for five months.

Table 57 Characteristics for creating a schedule to snapshot a volume

Item	Description and requirements
Name	The name of the snapshot created by 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” on page 34 for information about this naming convention. 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.
Description	[Optional] Must be from 0 to 127 characters.
Start at	The date and time can occur in the past.
Retention	The retention criteria can be for a specified number of snapshots or for a designated period of time.

Currently, you cannot create scheduled, application-managed snapshots from SAN/iQ. This function will be available in a future release.

1. In the navigation window, select the volume for which you want to create a schedule for snapshots.
The Volume tab window opens.
2. Click Volume Tasks on the Details tab and select New Schedule to Snapshot a Volume.
3. Type a name for the snapshots.
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. Specify the retention criteria for the snapshot.
9. Click OK when you have finished creating the schedule.
To view the schedule just created, select the Schedules tab view.

Editing scheduled snapshots

You can edit everything in the scheduled snapshot window except the name.

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. Click OK.

Pausing and resuming scheduled snapshots

At times it may be convenient to prevent a scheduled snapshot from taking place. Use these steps to pause and then resume a snapshot schedule.

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.

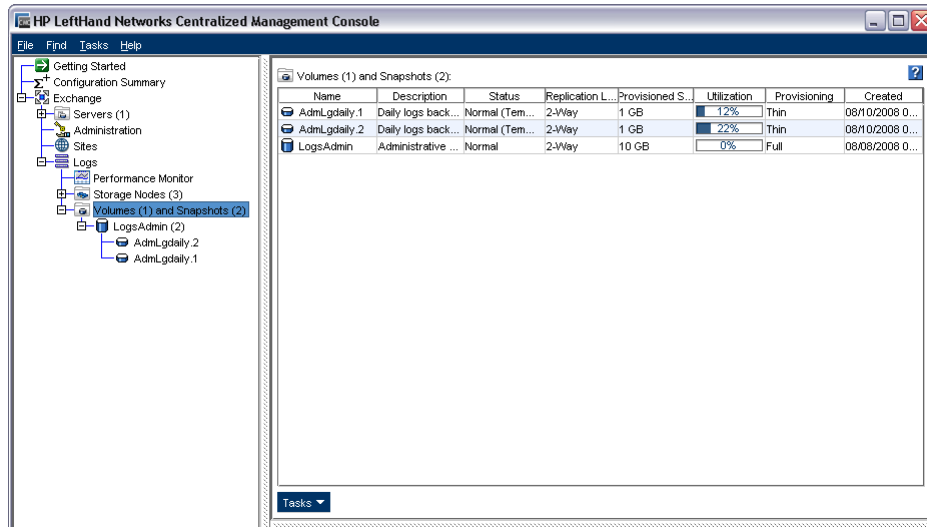


Figure 114 Delete multiple snapshots from the volumes and snapshots node

Scripting snapshots

Application-based scripting is available for taking snapshots. Using application-based scripts allows automatic snapshots of a volume. For detailed information, see [Chapter 16](#) on page 287 and the *Cliq User Manual*, found in the Documentation directory under the CMC Program Files, for information about the SAN/iQ command line interface.

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?](#)” on page 263.
- Use Remote Copy to copy the newer snapshots that you want to keep, before performing the rollback. See the *Remote Copy User Manual* for more information about copying data.

New in release 8.0

When rolling back a volume to a snapshot, the volume retains the original name. Releases before 8.0 required a new name for the rolled back volume.

Requirements for rolling back a volume

Best Practices

- 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. A future release will identify snapshot sets in the CMC. For more information, see [“Creating snapshots for volume sets”](#) on page 249.

Prerequisite

- If you need to preserve the original volume, or any snapshots that are newer than the one you will use for rolling back, use Remote Copy to create a copy of the volume or snapshots before beginning the roll back operation.

△ CAUTION:

When performing a roll back, 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.

Rolling back a volume from 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.

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. This message illustrates all the possible consequences of performing a roll back, 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.

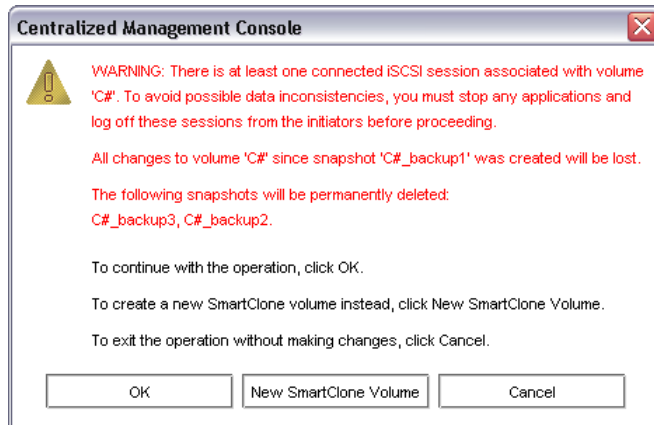


Figure 115 Rolling back a volume

Choosing a roll back strategy

You have three choices for continuing from this message window.

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.

1. Click OK to continue.
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.
2. If you rolled back an application-managed snapshot, use `diskpart.exe` to change the resulting volume's attributes.
For more information, see [“Making an application-managed snapshot available”](#) on page 250.
3. 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.

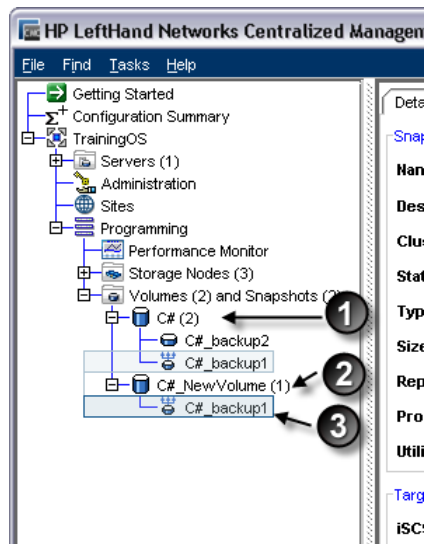
2. Enter a name and configure the additional settings.

For more information about characteristics of SmartClone volumes, see [“Defining SmartClone volume characteristics”](#) on page 267.

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.



1. Original volume
2. New SmartClone volume from snapshot
3. Shared clone point

Figure 116 New volume with 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 an application-managed snapshot available”](#) on page 250.

Cancel the roll back 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 roll back.

1. Click Cancel.
2. Perform necessary actions.
3. Start the roll back 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.

△ CAUTION:

Typically, you do not want to delete individual snapshots that are part of a snapshot set. A future release will identify snapshot sets in the CMC. For information about snapshot sets, see [“Requirements for application-managed snapshots”](#) on page 248. Typically, you want 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.

Prerequisites

- Stop any applications that are accessing the snapshot and log off all related iSCSI sessions

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.
A confirmation message opens.
5. Click OK.

15 SmartClone volumes

Overview of SmartClone volumes

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. 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](#)” on page 228.

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.

Glossary

Table 58 lists terms and definitions used for the SmartClone volumes feature. The illustration in Figure 117 on page 264 shows how the SmartClone volumes and related elements look in the CMC.

Table 58 Terms used for SmartClone features

Term	Definition
SmartClone Volume	A volume created using the SmartClone process. In Figure 117, the volume C#class_1 is a SmartClone volume.
Clone point 	The snapshot from which the SmartClone volumes are created. The clone point cannot be deleted. In Figure 117, the snapshot C#_SCsnap 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 117, the snapshots C#_snap1 and C#_snap2 are shared snapshots.
Map view	Tab that displays the relationships between clone points and SmartClone volumes. See the map view in Figure 131 on page 280 and Figure 132 on page 280.

In Figure 117 you can see a regular volume with 3 snapshots on the left and on the right, a regular volume with 1 SmartClone volume, 1 clone point and 2 shared snapshots.

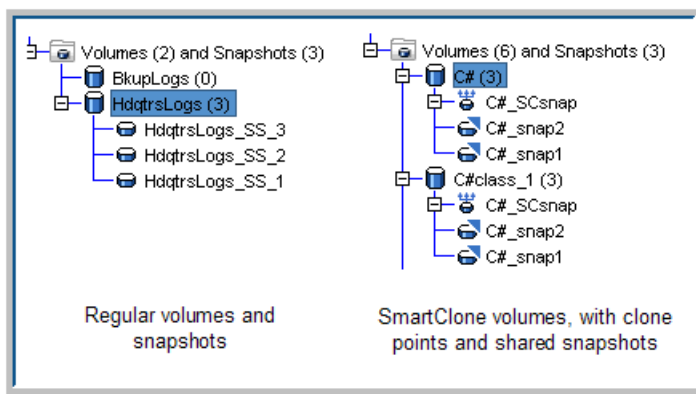


Figure 117 How SmartClone volumes, clone points and shared snapshots appear in the CMC

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 LeftHand Storage Solution volume.

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 LeftHand Storage 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

Say you want to track monthly trends in web requests for certain types of information. Monthly 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 replication 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 [Chapter 12](#) on page 221 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 LeftHand 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.
- Replication level - The replication level of the source volume must be retained when creating SmartClone volumes, though you can change the replication level after the SmartClone volumes are created. However, if you change the replication level for any SmartClone volume, the replication 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 resignaturing be enabled and will automatically name duplicate datastores. Most servers allow the duplicate disks to be renamed.

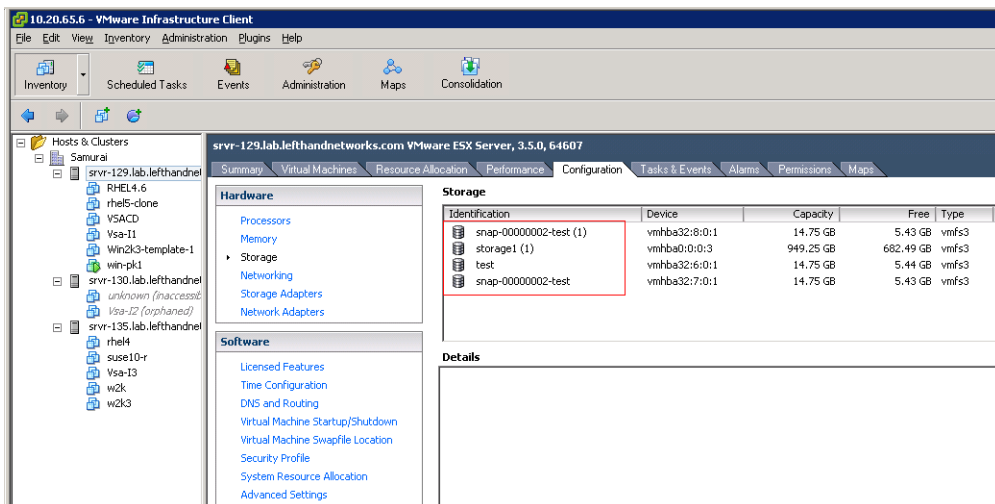


Figure 118 Duplicate names on duplicate datastores in ESX Server

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 [Chapter 17](#) on page 289.

Defining SmartClone volume characteristics

When creating SmartClone volumes, you define the following characteristics.

Table 59 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 more information, see “Assigning server connections access to volumes” on page 292.
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”](#) on page 34.

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 119](#) shows a SmartClone volume with the base name of “C#” and 10 clones. (The number in parenthesis indicates how many snapshots are under that volume.)

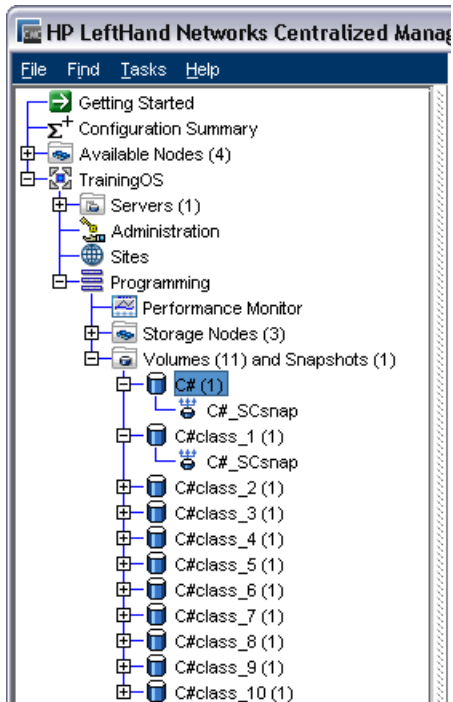
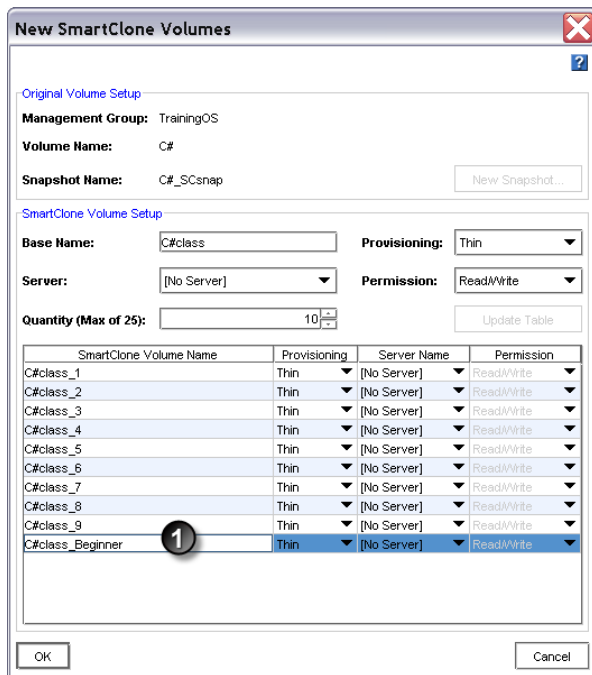


Figure 119 Example of using a base name with 10 SmartClone volumes

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 won't be disrupted.



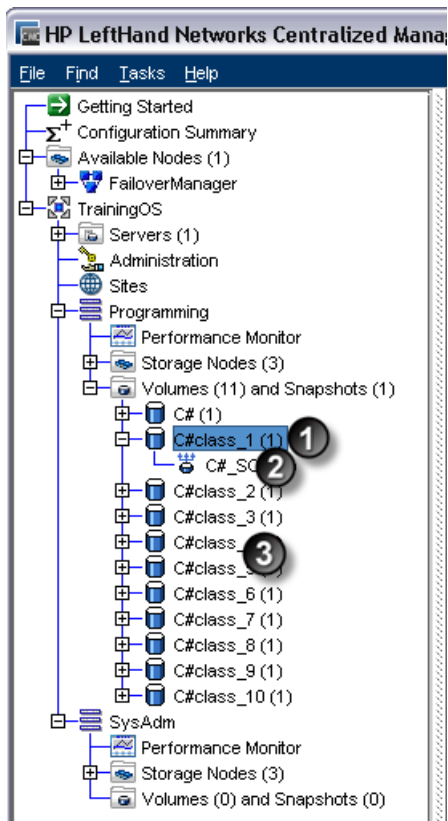
1. Rename SmartClone volume in list

Figure 120 Rename SmartClone volume from base name

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 associated volumes, which are noted in the message.

For example, in [Figure 121](#), 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, C#class_1, to the cluster SysAdm.



1. Source volume
2. Clone point
3. SmartClone volumes (10)

Figure 121 Programming cluster with 10 SmartClone volumes, 1 clone point, and the source volume

So you edit the volume `C#class_1` and on the Advanced tab you change the cluster to `SysAdm`. A confirmation message window opens. This message lists all the volumes and snapshots that will have their cluster changed as a result of changing `C#class_1`. In this case there are 12 volumes and snapshots that will move to cluster `SysAdm`, the original `C#` volume and the 10 SmartClone volumes, plus the clone point.

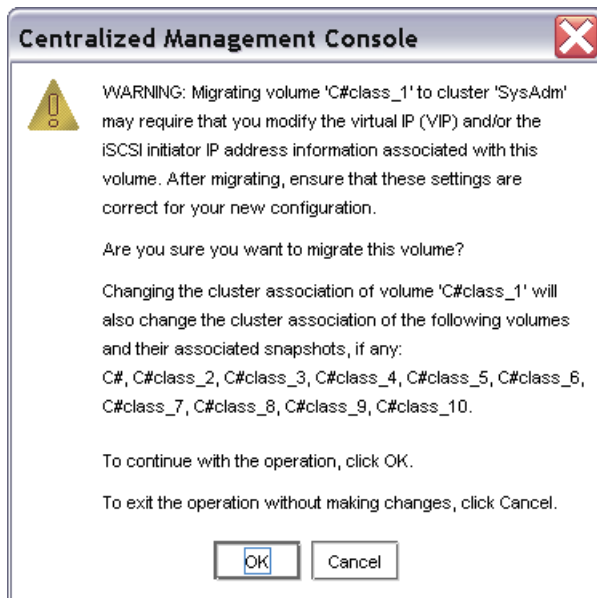


Figure 122 Changing one SmartClone volume changes all associated volumes and snapshots

When you click OK on the message, the 12 volumes and snapshots move to the cluster SysAdm.

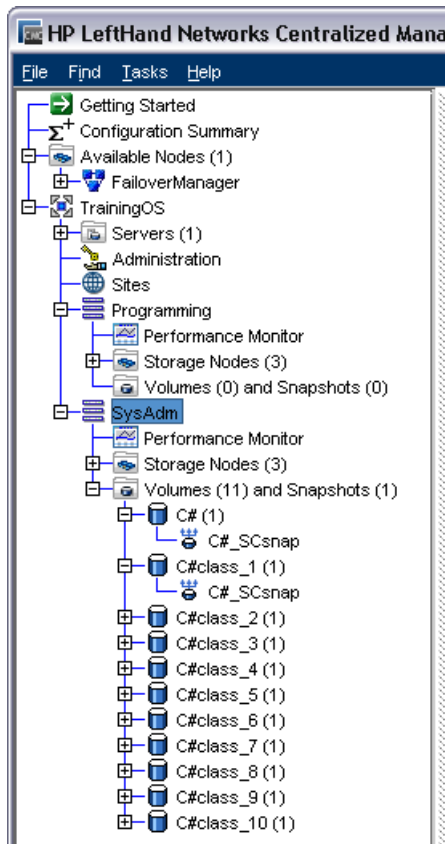


Figure 123 SysAdm cluster now has the 10 SmartClone volumes, 1 clone point, and the source volume

Table 60 shows the shared and individual characteristics of SmartClone volumes. Note that if you change the cluster or the replication level of one SmartClone volume, the cluster and replication level of all the related volumes and snapshots will change.

Table 60 Characteristics of SmartClone volumes

Shared characteristics	Individual characteristics
Cluster	Name
Replication level	Description
Replication priority	Size
	Type (Primary or Remote)
	Provisioning (Thin or Full)
	Server



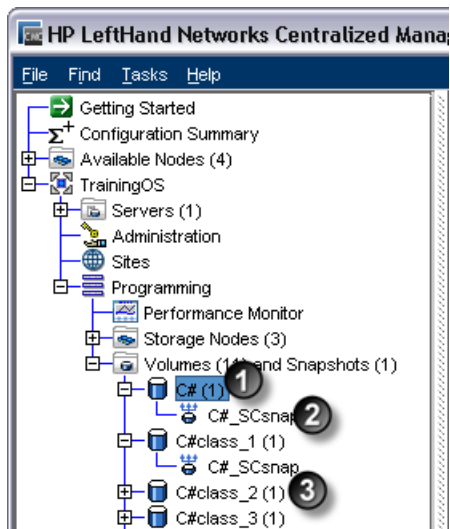
NOTE:

Snapshot schedules and remote copy schedules are also individual to a single SmartClone volume.

Clone point



The icon shown above 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.



1. Original volume
2. Clone point
3. SmartClone volume

Figure 124 Navigation window with clone point

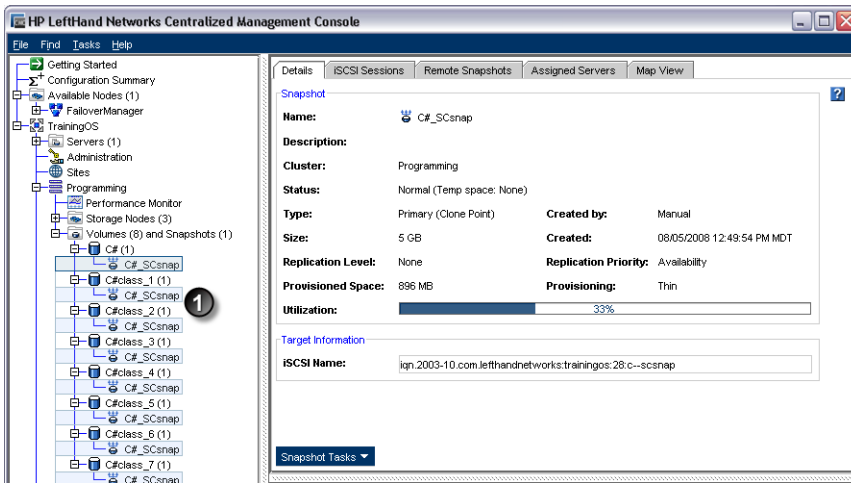
In [Figure 124](#), the original volume is “C#.”

- 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 61 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.



1. Clone point appears multiple times. Note that it is exactly the same in each spot

Figure 125 Clone point appears under each SmartClone volume



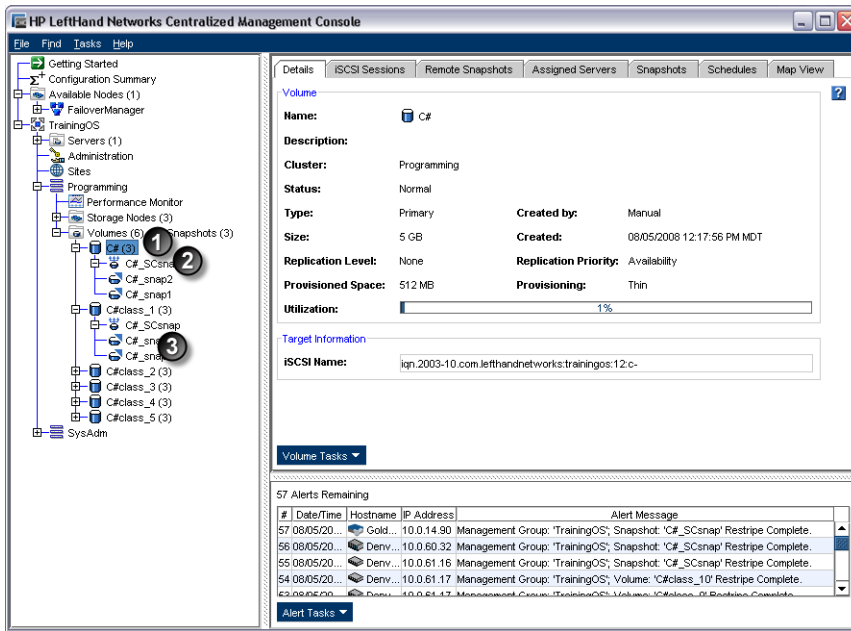
NOTE:

Remember! 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 above.



1. Original volume
2. Clone point
3. Shared snapshots



Figure 126 Navigation window with shared snapshots

In [Figure 126](#), 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 62 How it works - shared snapshots

First, a volume	C#
Next, 3 snapshots	C#_snap1 C#_snap2 C#_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 display under all the volumes which share them. In [Figure 126](#) on page 275, 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 set the characteristics for the entire group, or set them individually.



1. Set characteristics for multiples here
2. Edit individual clones here

Figure 127 Setting characteristics for SmartClone volumes

For details about the characteristics of SmartClone volumes, see “[Defining SmartClone volume characteristics](#)” on page 267.

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 a snapshot](#)” on page 248.

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.

- Click Update Table to populate the table with the number of SmartClone volumes you selected.

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#_SCsnap_1	Thin	[No Server]	Read/Write

- Enter # desired in Quantity field, and click Update Table

Figure 128 Creating multiple 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): 5 **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

OK Cancel

- List of SmartClone volumes created after clicking Update Table

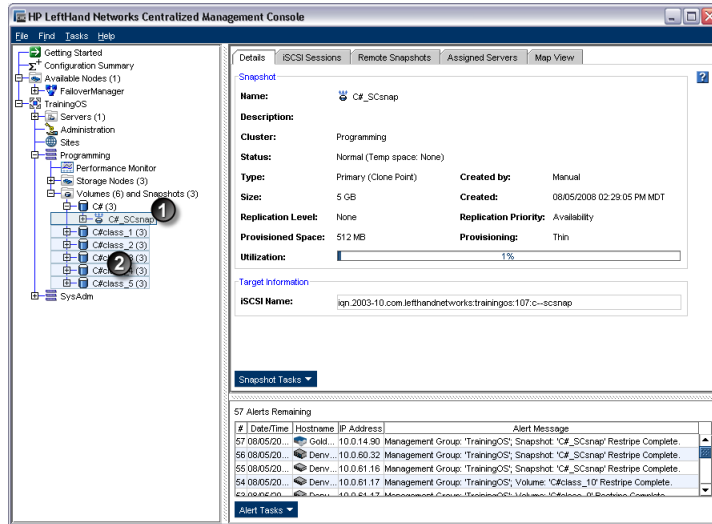
Figure 129 Creating multiple SmartClone volumes

- 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.



1. Clone point
2. New SmartClone volumes

Figure 130 New SmartClone volumes in Navigation window

Viewing SmartClone volumes

As you create multiple SmartClone volumes, you can view them and their associated volumes and snapshots in both the navigation window and in the Map View tab, shown in Figure 131 on page 280.

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 124 on page 273.

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.

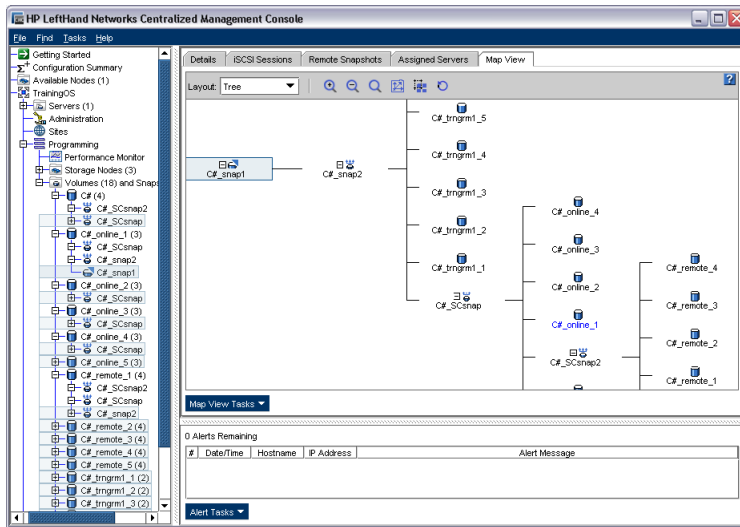


Figure 131 Viewing SmartClone volumes and snapshots as a tree in the Map View

Using views

The default view is the tree layout, displayed in Figure 131. 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.

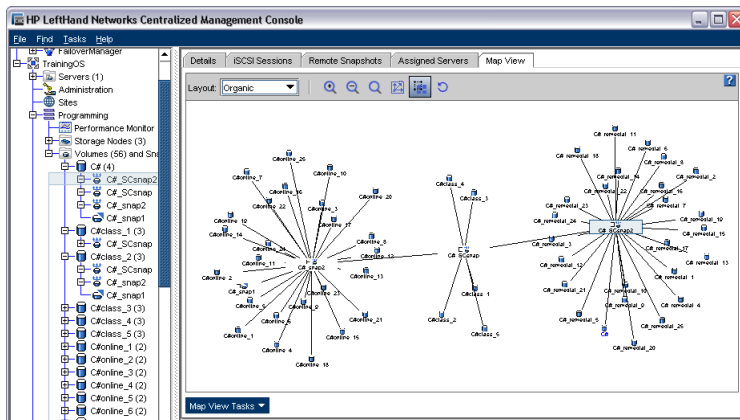


Figure 132 Viewing the organic layout of SmartClone volumes and associated snapshots in the Map View

Manipulating the Map View

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

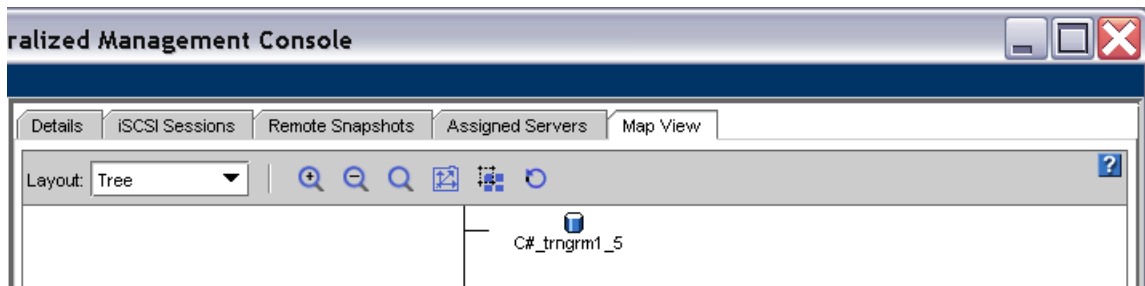








Figure 133 Toolbar with display tools in the Map View window

Using the display tools

Use these tools, described in Table 63, to select specific areas of the map to view, zoom in on, rotate, and move around the window. If you have a complex configuration of SmartClone volumes, use the Map View tools to easily view and monitor the configuration. Figure 134 (page 281) shows an example of the Magnify tool.

Table 63 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

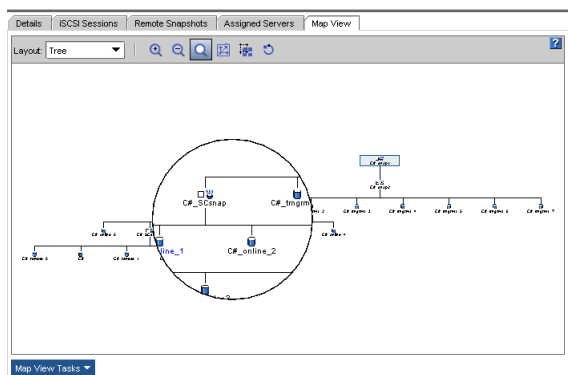
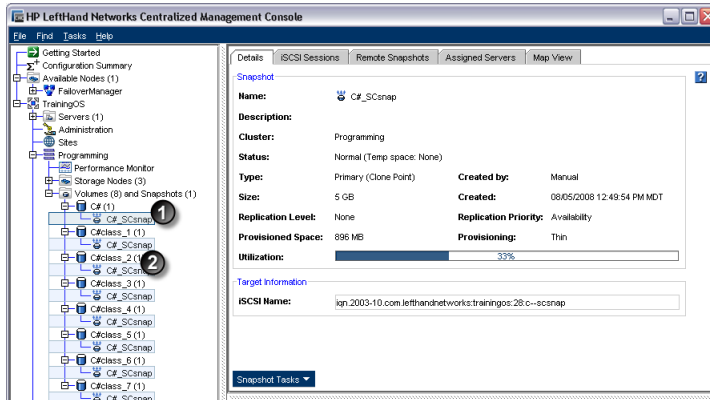


Figure 134 Using the Magnify tool with Map View tree

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 135](#), the clone point is selected in the tree. The clone point supports the 7 C# training class SmartClone volumes so it is displayed under those 7 volumes. The highlight shows the relationship of the clone point to the original volume plus the 7 SmartClone volumes created from the original volume.



1. Selected clone point
2. Clone point repeated under SmartClone volumes

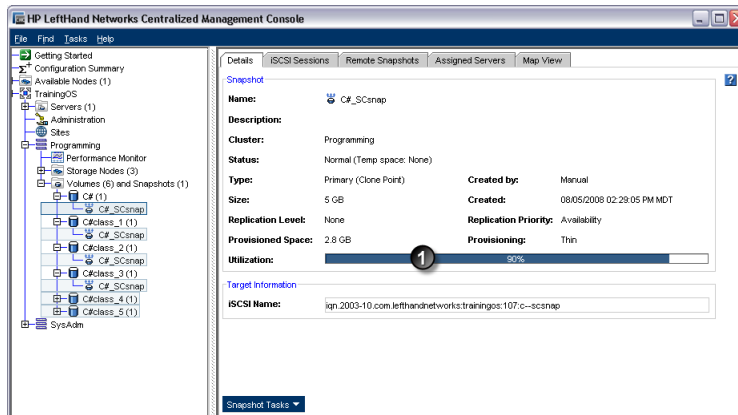
Figure 135 Highlighting all related clone points in navigation window

Viewing utilization of clone points and SmartClone volumes

Multiple SmartClone volumes share data from the clone point without requiring that data be duplicated for each SmartClone volume. On the Details tab of the clone point and the SmartClone volumes, there is a Utilization graph. Compare the Utilization graph for the clone point, and then for the SmartClone volumes. Note that the clone point contains data that is shared by SmartClone volumes, and the volumes themselves do not contain separate copies of that data. That is why the volumes' Utilization graphs show 0%.

In the example below, the clone point is at 90% of its 5 GB capacity with the C# training class desktop configuration. The 5 SmartClone volumes shared out for the 5 individual users contain no data at the time they are created for use by those 5 individuals. Only as each user writes data to his or her individual volume through the file share system mounted on that volume, do those volumes fill up on the SAN.

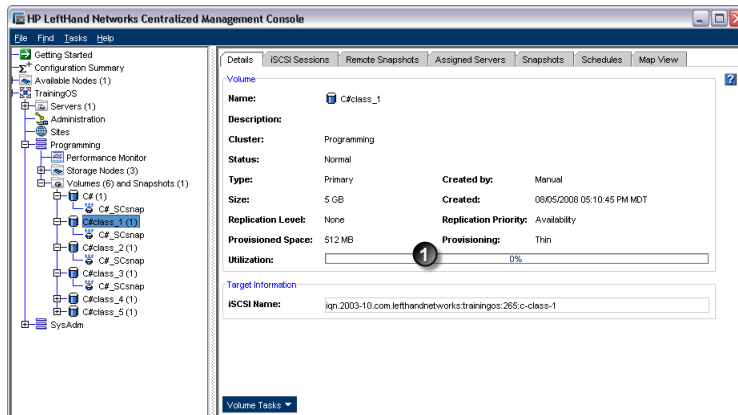
[Figure 136](#) shows the utilization graph of the clone point on the Details tab.



1. Utilization graph for clone point

Figure 136 Clone point Details tab showing utilization graph

Figure 137 shows the utilization of the SmartClone volume created from the clone point.



1. Utilization graph at 0%

Figure 137 SmartClone volume Details tab showing utilization graph

Editing SmartClone volumes

Use the Edit Volume window to change the characteristics of a SmartClone volume.

Table 64 Requirements for changing SmartClone volume characteristics

Item	Shared or Individual	Requirements for Changing
Description	Individual	May be up to 127 characters.
Size	Individual	Available space on cluster.
Servers	Individual	Existing server defined.

Item	Shared or Individual	Requirements for Changing
Cluster	Shared	<p>All associated 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 nodes and unallocated space for the size and replication level of the volume and all the other associated volumes and snapshots being moved. <p>When moving volumes to a different cluster, those volumes temporarily exist on both clusters.</p>
Replication Level	Shared	All associated volumes and snapshots must change to the same replication level. The cluster must have sufficient storage nodes and unallocated space to support the new replication level for all related volumes.
Replication Priority	Shared	All associated volumes and snapshots must change to the same replication priority. To change the replication priority, the replication level must support the change. You can always go from Redundancy to Availability. However, you cannot go from Availability to Redundancy unless a sufficient number of storage nodes in the cluster are available. For a detailed explanation, see Table 53 (page 240). For example, if you have 2-way replication with 3 storage nodes in the cluster, you can change from Availability to Redundancy if all the storage nodes in the cluster are available. For a detailed explanation about, see Replication priority (page 225).
Type	Individual	Whether the volume is primary or remote.
Provisioning	Individual	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.

The Edit Volume window opens. See [Table 64](#) on page 283 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 associated 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 such time that 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 depend on that clone point. After you delete all but one volume that depend 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 138](#), you must delete any four of the five C#class_x volumes before you can delete the clone point.

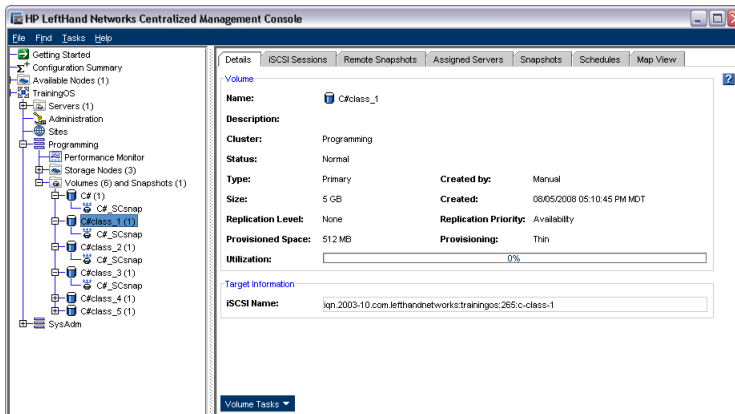


Figure 138 Viewing volumes that depend on a clone point

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.

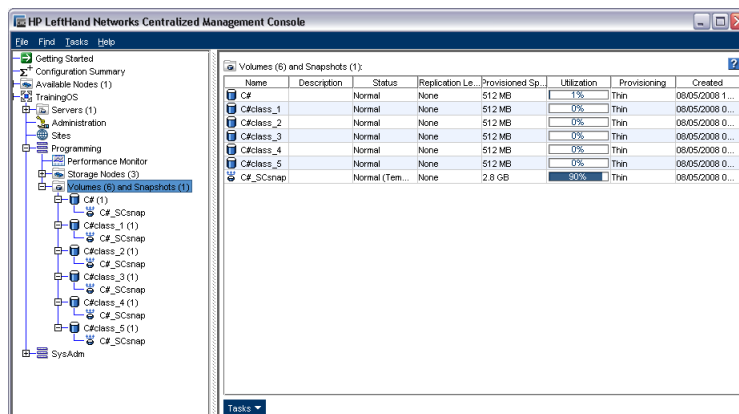


Figure 139 List of SmartClone volumes in cluster

2. Use Shift+Click to select the SmartClone volumes to delete.

3. Right-click and select Delete Volumes and Snapshots.

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.

16 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 is replaced by SAN/iQ CLIQ, the HP LeftHand Storage Solution command-line interface (CLI). 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 Command-Line Interface User Manual is available from the HP LeftHand Networks 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 LeftHand Networks website.

17 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 SAN/iQ software. We refer to this setup as a “server connection.”

You can set up servers to connect to volumes 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 SAN/iQ software to balance connections to the SAN.
- HP LeftHand 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 [Chapter 22](#) on page 335.

Setting up server connections to volumes requires the general tasks outlined below.

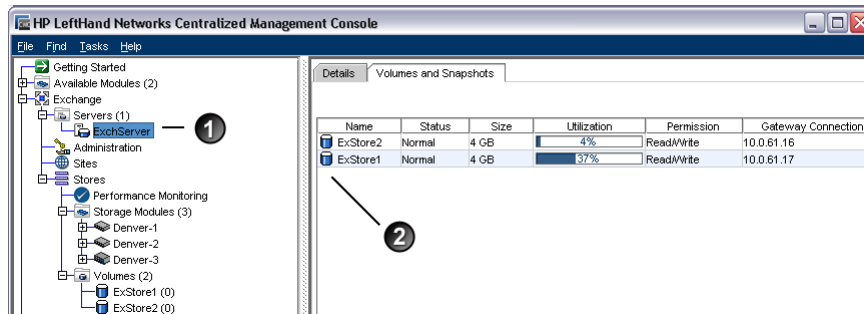
Table 65 Overview of configuring server access to volumes

Do This	For More Information
1. Ensure that an iSCSI initiator is installed on the server.	If you are using HP LeftHand DSM for MPIO, ensure that the Microsoft MPIO and the SAN/iQ HP LeftHand DSM for MPIO are installed on the server also. Refer to the HP LeftHand P4000 Windows Solution Pack User Manual.
2. 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 ” on page 290.
3. In the CMC, assign volumes to the server.	See “ Assigning server connections access to volumes ” on page 292.
4. In the iSCSI initiator on the server, log on to the volume.	See “ Completing the iSCSI Initiator and disk setup ” on page 294.
5. On the server, configure the volume using disk management tools.	See “ Completing the iSCSI Initiator and disk setup ” on page 294.

Former terminology (release 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 work 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.



1. Server connection in the navigation window with one server
2. Volumes and Snapshots tab shows two assigned volumes that the server can access

Figure 140 Server assignments in the navigation window and the Volumes and Snapshots tab

Adding server connections to management groups

Add each server connection that needs access to a volume to the management group where the volume exists. 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”](#) on page 292.

Prerequisites

- Each server must have an iSCSI initiator installed.
 - You know where to find the initiator node name in the iSCSI initiator. See [“iSCSI and CHAP terminology”](#) on page 338.
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.
The server connection name is case sensitive. It cannot be changed later, unless you delete and recreate the connection.
 4. Select the check box to allow access via iSCSI.
 5. If you plan to use iSCSI load balancing, click the link in the window to review the list of compliant iSCSI initiators.

Scroll down to see the entire list. If your initiator is not on the list, do not enable load balancing. For more information about iSCSI load balancing, see [“iSCSI load balancing”](#) on page 336.

△ CAUTION:

Using a non-compliant iSCSI initiator with load balancing can compromise volume availability during iSCSI failover events.

6. If you want to use iSCSI load balancing and your initiator is compliant, select the check box to enable load balancing.

7. Select the authentication method, either CHAP not required or CHAP required.
For more information, see [“Authentication \(CHAP\)”](#) on page 337.
8. In the Initiator Node Name field, enter the iqn string.
Open your iSCSI initiator and look for the string there. You can copy the string and paste it into the field.
For more information, see [“iSCSI and CHAP terminology”](#) on page 338.
9. If you are using CHAP, complete the fields necessary for the type of CHAP you intend to configure, as shown in [Table 66](#).

Table 66 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

10. Click OK.
The server connection displays 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”](#) on page 292.

Editing server connections

You can edit the following fields for a server connection:

- Description
- Load balancing
- CHAP options

You can also delete a server connection from the management group. For more information, see [“Deleting server connections”](#) on page 292.

△ CAUTION:

Editing a server 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.

If you change the Enable Load Balancing option, a warning message opens when you finish filling out this window. After changing the iSCSI load balancing configuration, you have to log your servers off, then log them back on to the volumes.

△ **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.

5. Click OK when you are finished.
6. If you have changed the Enable Load Balancing option, you must log servers off the volumes. This may entail stopping the applications, disconnecting them, reconnecting the applications to the volumes, and then restarting them.

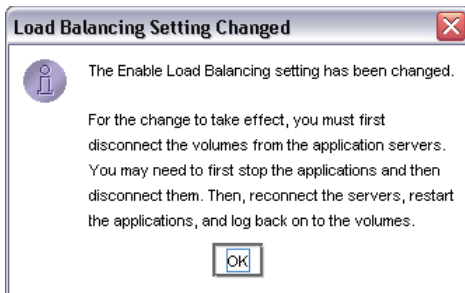


Figure 141 Warning after changing load balancing check box

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.

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 two ways:

- [“Assigning server connections from a volume”](#) on page 293
- [“Assigning volumes from a server connection”](#) on page 293

Prerequisites

- The server connections you want to assign must already exist in the management group. See [“Adding server connections to management groups”](#) on page 290.
- The volumes or snapshots you want assign must already exist in the management group. See [“Creating a volume”](#) on page 239.

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 67](#).

Table 67 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”](#) on page 292.

1. In the navigation window, right-click the volume you want to assign server connections to.
2. Select Assign and Unassign Servers.
3. Click the Assigned check box for each server connection you want to assign to the volume or snapshot.
4. From the Permission drop-down list, select the permission each server connection should have to 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”](#) on page 294.

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”](#) on page 292.

1. In the navigation window, right-click the server connection you want to assign.
2. Select Assign and Unassign Volumes and Snapshots.

3. Click the Assigned check box for each volume or snapshot you want to assign to the server connection.
4. From the Permission drop-down 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](#)" on page 294.

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

Editing server connection assignments from a volume

You can edit the assignment of one or more server connections to any volume or snapshot.

△ 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.

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.

△ 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.

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 [Chapter 22](#) on page 335.

Refer to the operating system-specific documents in the Resource Center for more information about setting up volumes and iSCSI.

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 LeftHand DSM for MPIO settings

If you are using HP LeftHand 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 HP LeftHand DSM for MPIO, refer to the HP LeftHand P4000 Windows Solution Pack User Manual.

Disk management

You must also format, configure, and label the volumes from the server using the operating system’s disk management tools.

18 Monitoring performance

The Performance Monitor provides performance statistics for iSCSI and storage node I/Os to help you and HP LeftHand Networks 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 [Chapter 7](#) on page 129.

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 nodes?), and
- Data placement (Should this volume be on my SATA or SAS cluster?).

The performance data do 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 cluster, volume, and storage node 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 node performance data will allow you to easily isolate, for example, a specific storage node with higher latencies than the other storage nodes in the cluster.

Prerequisites

- You must have a cluster with one or more storage nodes and one or more volumes connected via iSCSI sessions.
- All storage nodes 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 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”](#) on page 309.

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 like crazy. What's happening?

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 142 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 143 Example showing volume's type of workload

Fault isolation example

This example shows that the Denver-1 storage node (dotted line pegged at the top of the graph) has a much higher IO read latency than the Denver-3 storage node. 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 144 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 145 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 146 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 147 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 nodes improve performance?
- Is the load between two clusters balanced? If not, what should I do?
- I have budget to purchase two new storage nodes.
 - 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 nodes 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 nodes. You can see that Denver-1 averages more than 79% utilization. You can increase the networking capacity available to that storage node by enabling NIC bonding on the storage node. You can also spread the load out across the storage nodes using iSCSI load balancing.




Figure 148 Example showing network utilization of three storage nodes

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.

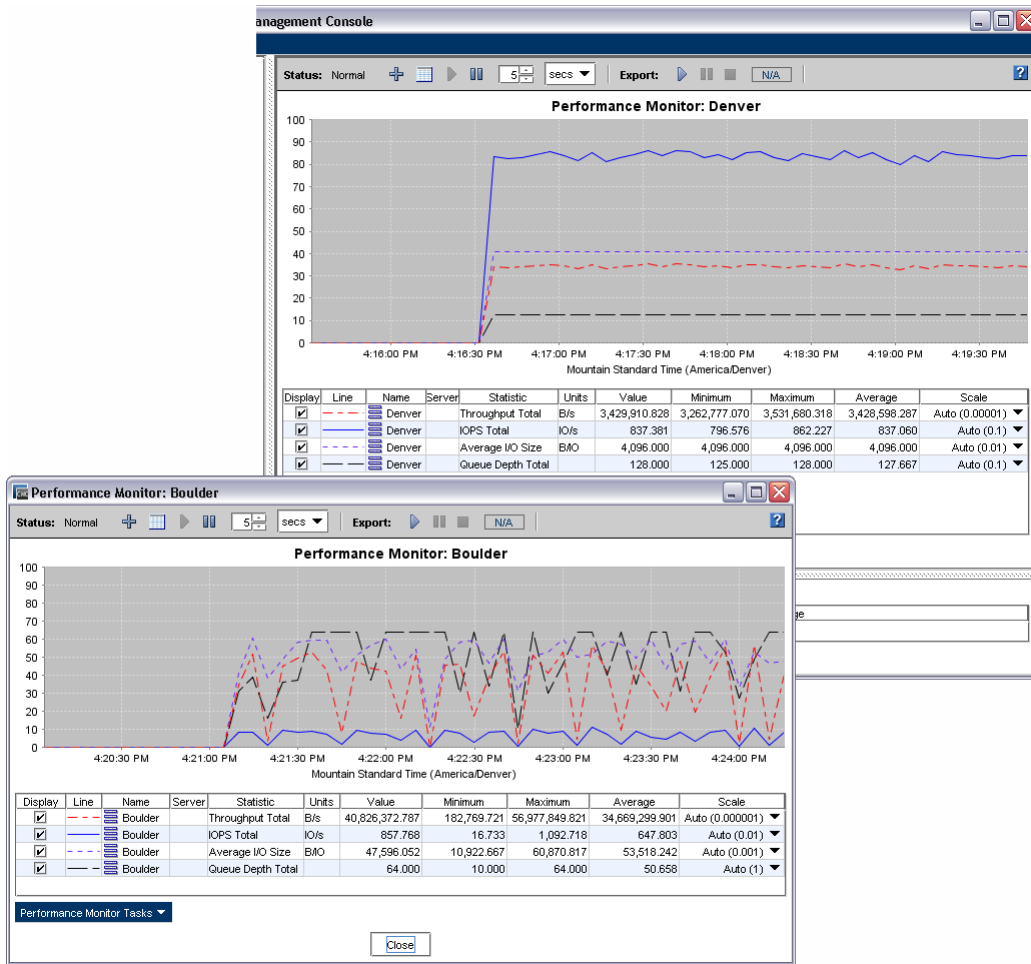


Figure 149 Example comparing two clusters

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 150 Example comparing two volumes

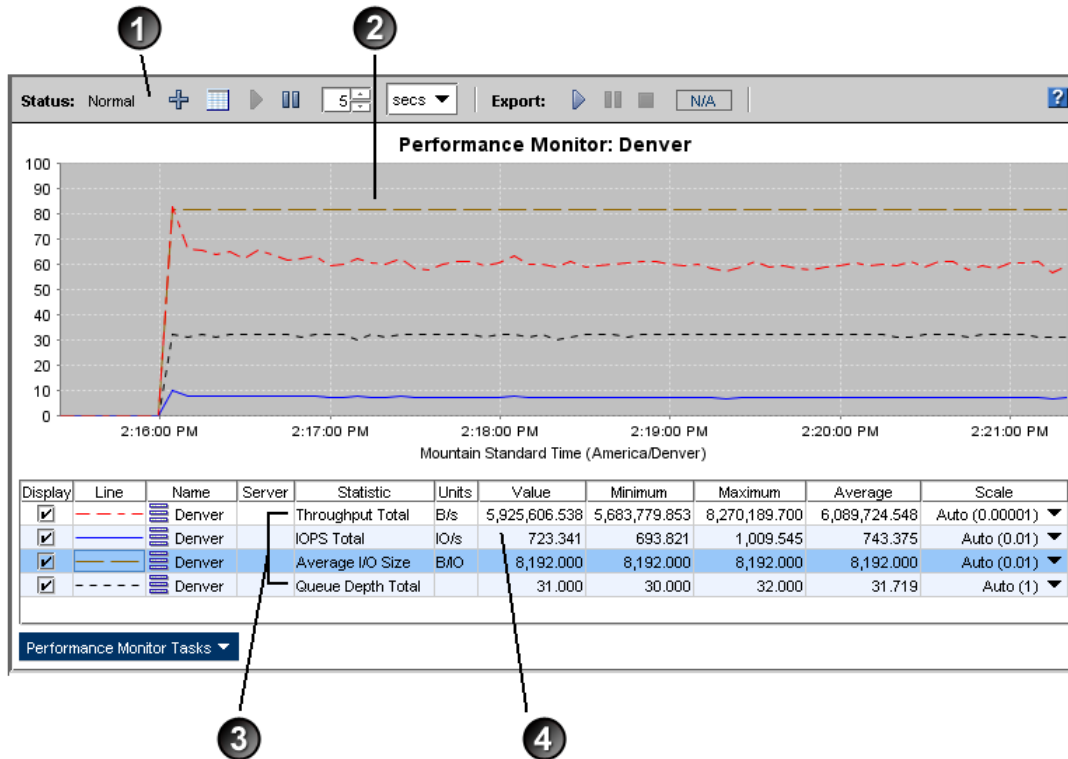
Accessing and understanding the Performance Monitor window

The Performance Monitor is available as a tree node 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 node 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.



1. Toolbar
2. Graph
3. Default statistics
4. Statistics table

Figure 151 Performance Monitor window and its parts

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” on page 304
- “Performance monitoring graph” on page 305

- “Performance monitoring table” on page 306

Performance Monitor toolbar

The toolbar lets you change some settings and export data.



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 nodes. Click the Warning text for more information. Figure 153 on page 305
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.

Figure 152 Performance Monitor toolbar

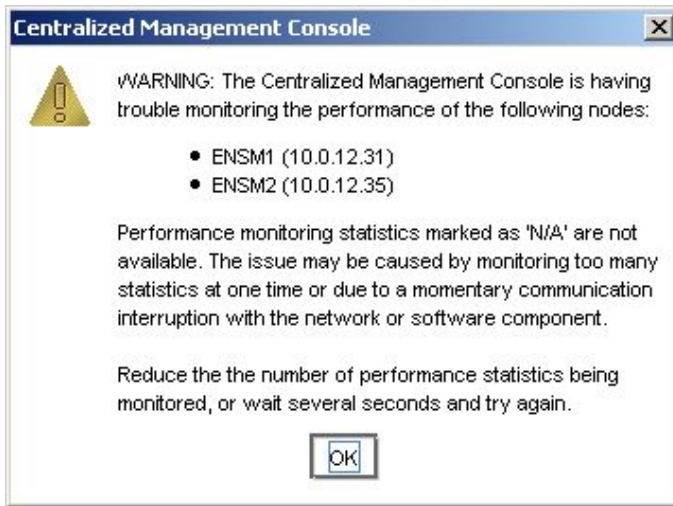


Figure 153 Example of a warning message

Performance monitoring graph

The performance monitor graph shows a color-coded line for each displayed statistic.




Figure 154 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”](#) on page 315.

The horizontal axis shows the either 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”](#) on page 309. (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”](#) on page 313.

Performance monitoring table

The performance monitor table displays a row for each selected statistic.

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...	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) ▼

Figure 155 Performance Monitor table

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 68](#) on page 306 defines each column of the Performance Monitor table.

Table 68 Performance Monitor table columns

Column	Definition
Display	Toggles 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 node, or volume being monitored.
Server	For volumes and snapshots, shows the server that has access.
Statistic	The statistic you selected for monitoring.
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"](#) on page 310.

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 nodes 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 nodes and devices are reporting is the overhead (replication, remote snapshots, and management functions).




Figure 156 Performance statistics and where they are measured

The following statistics are available:

Table 69 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

Statistic	Definition	Cluster	Volume or Snapshot	NSM
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 node for the sample interval.	-	-	X
Memory Utilization	Percent of total memory used on this storage node for the sample interval.	-	-	X
Network Utilization	Percent of bi-directional network capacity used on this network interface on this storage node 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 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.
The Performance Monitor window opens.
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.
The Performance Monitor window opens.
3. Click Performance Monitoring Tasks and select Edit Monitoring Interval.
The Edit Monitoring Interval window opens.
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”](#) on page 306.

1. In the navigation window, log in to the management group.
2. Select the Performance Monitor node for the cluster you want.
The Performance Monitor window opens.

3. Click .

The Add Statistics window opens.

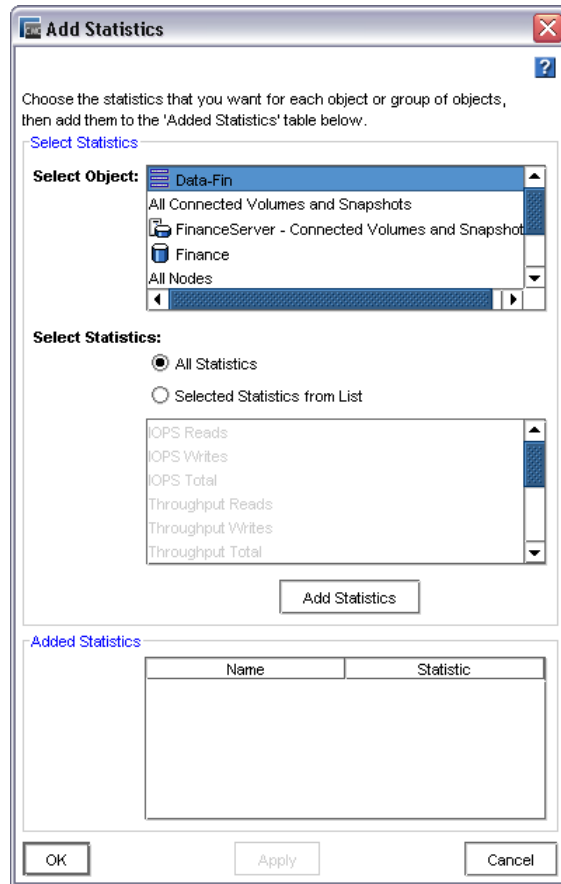


Figure 157 Add Statistics window

4. From the Select Object list, select the cluster, volumes, and storage nodes 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.

The statistics you selected are listed in the Added Statistics list.

If you selected a statistic that is already being monitored, a message displays 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.
The Performance Monitor window opens.
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 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.
The Performance Monitor window opens.
3. Right-click a row in the table and select Remove Statistics.
Use the CTRL key to select multiple statistics from the table.
A message displays confirming that you want to remove the selected statistics.
4. Click OK.

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.
The Performance Monitor window opens.
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.
The Performance Monitor window opens.
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.
The Performance Monitor window opens.
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 update when the next sample interval elapses. The graph will have a gap in the time.



Changing the graph

You can change graph and its lines in the following ways:

- [“Hiding and showing the graph”](#) on page 314
- [“Displaying or hiding a line”](#) on page 314
- [“Changing the color or style of a line ”](#) on page 314
- [“Highlighting a line”](#) on page 315
- [“Changing the scaling factor”](#) on page 315

Hiding and showing the graph

By default, the performance monitor graph displays 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.

The Edit Line window opens.

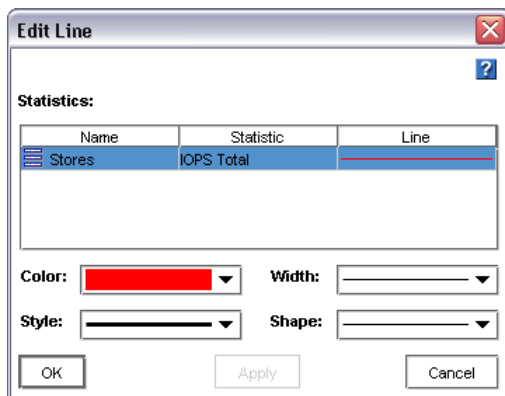


Figure 158 Edit Line window

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.

The Add Statistics window opens.

6. From the Select Object list, select cluster, volumes, and storage nodes 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 displays 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.
The Save window opens.
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.
You can save as a .png or .jpg.
6. Click Save.

19 Registering advanced features

Advanced features expand the capabilities of the SAN/iQ software. All advanced features are available when you begin using the SAN/iQ software. If you begin using a feature without first registering, a 30-day evaluation period begins. Throughout the evaluation period you receive reminders to register and purchase a license for the advanced features you want to continue using.

The advanced features are listed below.

- Multi-Node Virtualization and Clustering; clustered storage nodes 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.

Evaluating advanced features

Advanced features are active and available when you install and configure your system.

30—Day evaluation period

When you use any feature that requires registration, a message opens, asking you to verify that you want to enter a 30-day evaluation period.

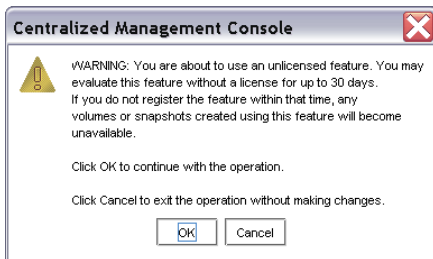


Figure 159 Verifying the start of the 30—day evaluation period

During this evaluation period you may configure, test, and modify any feature. At the end of the 30-day evaluation period, if you do not purchase a license key, then all volumes and snapshots associated with the feature become unavailable to any clients. The data is safe and you can continue to manage the volumes and snapshots in the CMC. Also, the entire configuration can be restored to availability when a license key is purchased and applied to the storage nodes in the management group that contains the configured advanced features.

 **NOTE:**

If you know you are not going to purchase the feature, plan to remove any volumes and snapshots created by using the feature before the end of the 30-day evaluation period.

Tracking the time remaining in the evaluation period

Track the time left on your 30-day evaluation period by using either the management group Registration tab or the reminder notices that open periodically.

Viewing licensing icons

You can check the status of licensing on individual advanced features by the icons displayed. Note that the violation icon displays throughout the 30-day evaluation period.




Figure 160 Icons indicating license status for advanced features

Starting the evaluation period

You start the evaluation period for an advanced feature when you configure that feature in the CMC.

Table 70 Descriptions of advanced features

Advanced Feature	Provides This Functionality	And Enters the License Evaluation Period When
Multi-Node Virtualization and Clustering	Clustering multiple storage nodes to create pools of storage.	You add 2 or more storage nodes 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.

Backing out of Remote Copy evaluation

If you decide not to purchase Remote Copy, 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 71](#), according to how you want to handle the data.

Table 71 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 30-day evaluation period without requiring that you use the CMC, you must first verify that you are aware of starting the 30-day 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.
 1. Read the text and check the box to enable the use of scripts during a license evaluation period.
 2. Click OK.

Turn off scripting evaluation

Turn off the scripting evaluation period when you take either one of these actions:

- You purchase the feature you were evaluating.
- You complete the evaluation and decide not to purchase any advanced features.

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 72 describes additional steps to safely back out of the scripting evaluation.

Table 72 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.

 **NOTE:**

Turning off the scripting evaluation ensures that no scripts continue to push the 30-day evaluation clock.

Registering advanced features

When registering storage nodes for advanced features, you must have your license entitlement certificate and submit the appropriate storage node feature key(s) to purchase the license key(s). You then receive the license key(s) and apply them to the storage node(s).

Using license keys

License keys are assigned to individual storage nodes. License keys can be added to storage nodes either when they are in the Available Nodes pool or after they are in a management group. One license key is issued per storage node and that key licenses all the advanced features requested for that storage node. Therefore, you register each storage node for which you want to use advanced features.

For example, if you wanted to configure 3 storage nodes in two locations to use with Remote Copy, you license the storage nodes in both the primary location and the remote location.

 **NOTE:**

If you remove a storage node from a management group, the license key remains with that storage node. See for more information about removing storage nodes from a management group.

Registering available storage nodes for license keys

Storage nodes that are in the Available Nodes pool are licensed individually. You license an individual storage node on the Feature Registration tab for that node.

The Feature Registration tab displays the following information:

- The storage node feature key, used to obtain a license key
- The license key for that storage node, if one has been purchased

- The license status of all the advanced features

Submitting storage node feature keys

1. In the navigation window, select the storage node from the Available Nodes 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. Go to <https://webware.hp.com> to register and generate the license key.

Entering license keys to storage nodes

When you receive the license keys, add them to the storage nodes.

1. In the navigation window, select the storage node from the Available Nodes 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.

- Click OK.

The license key appears in the Feature Registration window.

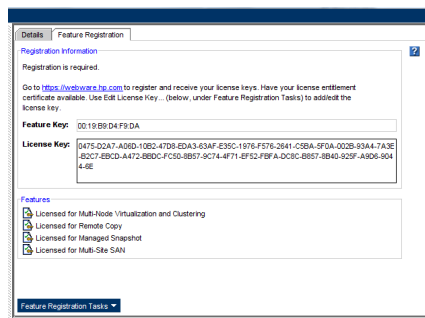


Figure 161 Storage node with a license key

Registering storage nodes in a management group

Storage nodes that are in a management group are licensed through the management group. You license the storage nodes 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 30-day evaluation period and which advanced features are in use and not licensed.

- Version information about software components of the operating system
- Customer information

Submitting storage node feature keys

Submit the feature keys of all the storage nodes in the management group.

1. In the navigation window, select the management group for which you want to register advanced features.
2. Select the Registration tab.

The Registration tab lists purchased licenses. If you are evaluating advanced features, the time remaining in the evaluation period is listed on the tab as well.

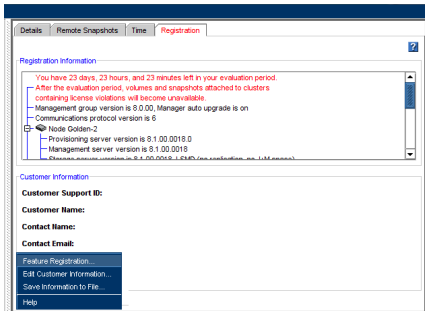


Figure 162 Registering advanced features for a management group

1. Click Registration Tasks and select Feature Registration from the menu.

The Feature Registration window lists all the storage nodes in that management group.

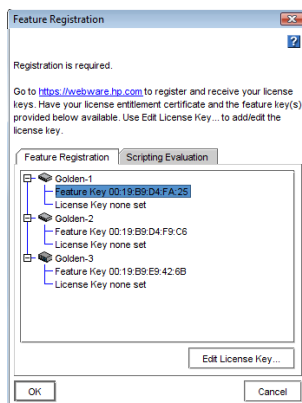


Figure 163 Selecting the feature key

2. For each storage node listed in the window, select the Feature Key.
3. Right-click and press Ctrl+C to copy the Feature Key.
4. Use Ctrl+V to paste the feature key into a text editing program, such as Notepad.
5. Go to <https://webware.hp.com> to register and generate the license key.



NOTE:

Record the host name or IP address of the storage node along with the feature key. This record will make it easier to add the license key to the correct storage node when you receive it.

Entering license keys

When you receive the license keys, add them to the storage nodes 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. Select a storage node and click Edit License Key.



Figure 164 Entering a license key

5. Copy and paste the appropriate license key for that storage node into the window.



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.

6. Click OK.

The license key appears in the Feature Registration window.




Figure 165 Viewing license keys

7. Click OK again to exit the Feature Registration window.

Saving license key information

For record-keeping, save the license information to a 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.

The license information is saved to a .txt file.

Saving and editing your customer information

This section explains how to save your customer profile, registrations, and licensing information. If you have this saved as a text file, and lose a storage node, it can help in the rebuild of a new storage node.

Make a customer information file for each management group in your system.

- First, create or edit your customer profile.
- Save the file to a computer that is not part of your storage system.

Editing your customer information file

Occasionally, you may want to change some of the information in your customer profile. If your company moves, or contact information changes for example.

1. In the navigation window, select a management group.
2. Click on the Registration tab to open that window.
1. Click Registration Tasks and select Edit Customer Information from the menu.
2. Fill in or change any of the information on this window.
3. 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.
1. Click Registration Tasks and select Save Information to File from the menu.
2. In the Save window, navigate to the directory where you want to save the license key and customer information file.
3. In the File Name field, enter a name for the file, which defaults to a .txt file.
4. Click Save.

Verify the information by viewing the saved .txt file.

20 SNMP MIB information

The SNMP Agent resides in the storage node. The agent takes SNMP network requests for reading or writing configuration information and translates them into internal system requests. Management Information Base (MIB) files are provided which can enable the system administrator to use their favorite SNMP tool to view or modify configuration information. The SNMP Agent supports versions 1 and 2c of the protocol.



NOTE:

To ensure that all items display properly in your SNMP tool, use version 2c or later of the protocol.

The supported MIBs

The following are the supported 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
- NOTIFICATION-LOG-MIB
- RFC1213-MIB
- SNMP-TARGET-MIB
- SNMP-VIEW-BASED-ACM-MIB
- SNMPv2-MIB
- SNMPv2-SMI
- UCD-DLMOD-MIB
- UCD-SNMP-MIB

21 Replacing disks appendix

This document 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 node. 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 nodes where RAID is running, but may be degraded, can be accomplished by following the procedures described in “[Replacing disks](#)” on page 328.

The following situations may require consulting with Customer Support to identify bad disks and then following the procedures below to rebuild the data (when replicated) on the storage node.

- RAID0 (Stripe) — RAID is off due to a failed disk.
- RAID5, 5+spare (Stripe with parity), and 50 — if multiple disks need to be replaced, then those disks must be identified and replaced, and the data on the entire storage node rebuilt.
- RAID10/1+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 node needs to be rebuilt.
- RAID6 (Stripe with dual parity) — if multiple disks need to be replaced, then those disks must be identified and replaced, and the data on the entire storage node rebuilt.

Before you begin

1. Know the name and physical location of the storage node that needs the disk replacement.
2. Know the physical position of the disk in the storage node.
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 replicated volumes and snapshots should show a status of Normal. Non-replicated volumes may be blinking.
- If volumes or snapshots are not replicated, change them to 2-way replication 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 nodes 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 for any one of the listed cases.

- RAID0 goes off.
- When multiple disks need to be replaced on a storage node with RAID5, RAID50, or RAID6.
- When multiple disks on the same mirror set need to be replaced on a storage node with RAID10.

Verify storage node not running a manager

Verify that the storage node that needs the disk replacement is not running a manager.

1. Log in to the management group.
2. Select the storage node in the navigation window and review the Details tab information. If the Storage Node 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 node in the navigation window and select Stop Manager.

When the process completes successfully, the manager is removed from the Status line in the Storage Node 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 the cluster has an odd number of managers, do one of these tasks:
 - Start a manager on another storage node or
 - Add a virtual manager to the management group by right-clicking on the management group name in the navigation window and select Add Virtual Manager.

Repair the storage node

Prerequisite

If there are non-replicated volumes that are blinking, you must either replicate them or delete them before you can proceed with this step. You see the message shown in [Figure 166](#) in this case.

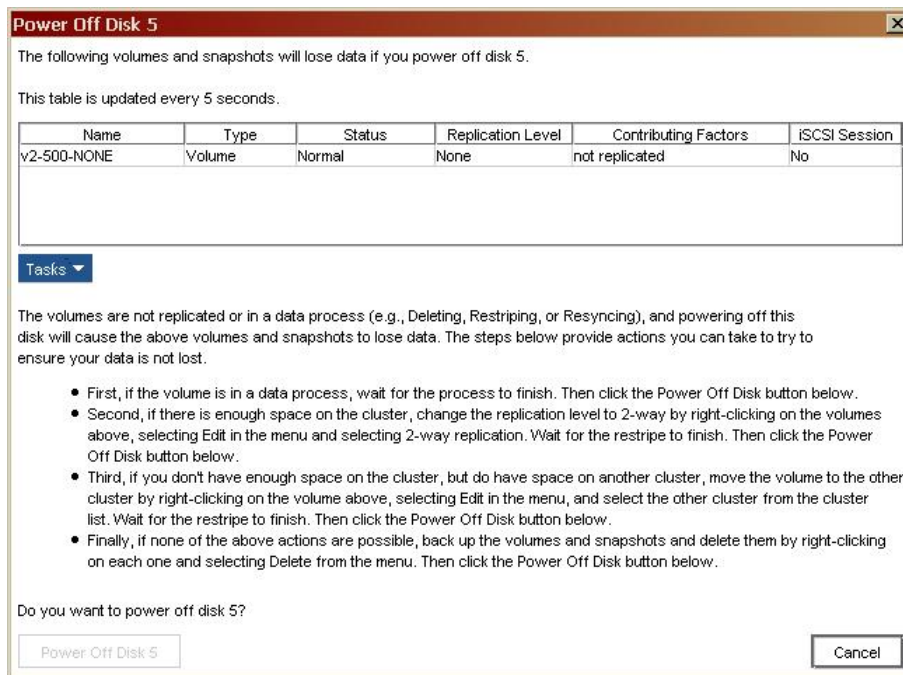


Figure 166 Warning if volumes are not replicated

- Right-click on the storage node in the navigation window and select Repair Storage Node. A “ghost” image takes the place of the storage node in the cluster with the IP address serving as a place holder. The storage node itself moves from the management group to the Available Nodes pool.

NOTE:

If the storage node does not appear in the Available Nodes pool, use the “Find” menu option to re-locate it.

Replace the disk

In the NSM 160 or the NSM 260

When these platforms are configured in RAID0, the menu choices for powering on and powering off disks are enabled.

1. Reconfigure the storage node for RAID0 if it is not already in RAID0.
2. In the CMC, power off the disk (you may power off up to 3 disks), one at a time.
See the procedures for powering off a disk in RAID0 in “[Manually power off the disk in the CMC for RAID0](#)” on page 83
3. Physically replace the disks in the storage node.
4. In the CMC, power each disk back on, one at a time.
After the last disk is powered on, RAID becomes Normal.

In the DL380 or the IBM x3650

△ CAUTION:

You must always use a new drive when replacing a disk in an IBM x3650. Never reinsert the same drive and power it on again.

1. Reconfigure the storage node for RAID0 if it is not already in RAID0.
2. Power off 1 disk.
See the procedures for powering off a disk in RAID0 in [“Manually power off the disk in the CMC for RAID0”](#) on page 83.
3. Physically replace the disk in the storage node.
4. Power on the disk.
5. Repeat [Step 2](#) through [Step 4](#) until all necessary disks are replaced.

In the DL320s (NSM 2120), Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300

For the DL320s (NSM 2120), Dell 2950, NSM 2060, NSM 4150, and HP LeftHand P4500, use the disk replacement procedures in [“Replacing a disk in a hot-swap platform \(NSM 160, NSM 260, DL380, DL320s \[NSM 2120\], Dell 2950, NSM 2060, NSM 4150, HP LeftHand P4500, HP LeftHand P4300\)”](#) on page 86.

△ CAUTION:

You must always use a new drive when replacing a disk in an Dell 2950, NSM 2060, or NSM 4150. Never reinsert the same drive or another drive from the same Dell 2950, NSM 2060, or NSM 4150.

Rebuilding data

The following steps take you through the steps to first rebuild RAID on the storage node and then to rebuild data on the storage node after it is added to the management group and cluster.

Re-create the RAID array

1. Select the Storage category and then 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 node and try reconfiguring the RAID again. If this second attempt is not successful, call customer support.

Checking progress for RAID array to rebuild

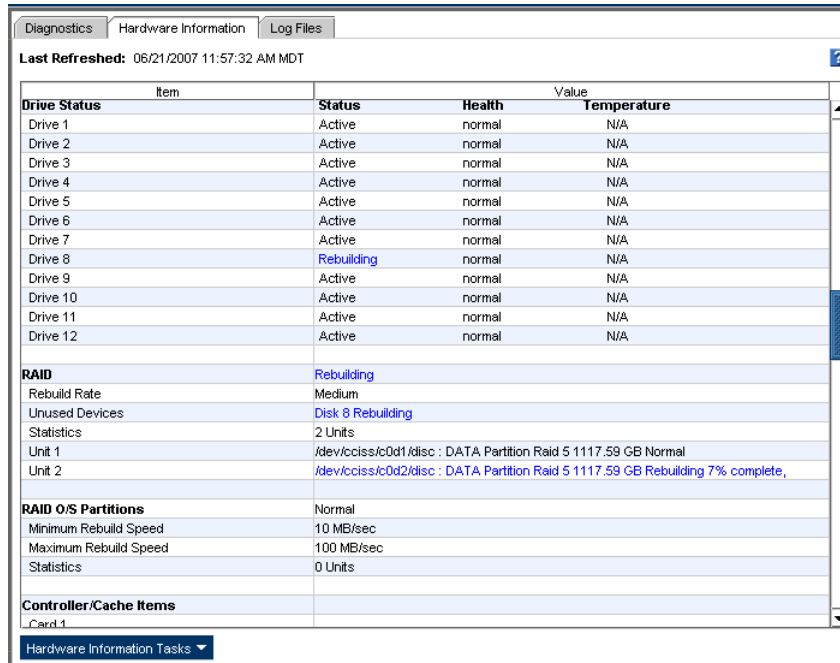
For NSM 160, NSM 260, DL380, DL320s (NSM 2120), IBM x3650, HP LeftHand P4500, HP LeftHand P4300 only

Use the Hardware Information report to check the status of the RAID rebuild.

1. Select the Hardware category and then select the Hardware Information tab.
2. Click the link on the tab "Click to Refresh" and scroll down to the RAID section of the Hardware report, shown in [Figure 167](#).

You can view the RAID rebuild rate and percent complete.

3. Click Hardware Information Tasks and select Refresh to monitor the ongoing progress.



Item	Status	Health	Temperature
Drive 1	Active	normal	N/A
Drive 2	Active	normal	N/A
Drive 3	Active	normal	N/A
Drive 4	Active	normal	N/A
Drive 5	Active	normal	N/A
Drive 6	Active	normal	N/A
Drive 7	Active	normal	N/A
Drive 8	Rebuilding	normal	N/A
Drive 9	Active	normal	N/A
Drive 10	Active	normal	N/A
Drive 11	Active	normal	N/A
Drive 12	Active	normal	N/A
RAID	Rebuilding		
Rebuild Rate	Medium		
Unused Devices	Disk 8 Rebuilding		
Statistics	2 Units		
Unit 1	/dev/cciss/c0d1/kdisc : DATA Partition Raid 5 1117.59 GB Normal		
Unit 2	/dev/cciss/c0d2/kdisc : DATA Partition Raid 5 1117.59 GB Rebuilding 7% complete,		
RAID O/S Partitions	Normal		
Minimum Rebuild Speed	10 MB/sec		
Maximum Rebuild Speed	100 MB/sec		
Statistics	0 Units		
Controller/Cache Items			
Card 1			

Figure 167 Checking RAID rebuild status

For Dell 2950, NSM 2060, and NSM 4150 only

Use the RAID Setup tab to check the status of the RAID rebuild.

1. Select the Storage category and then select the RAID Setup tab.
2. Click the link on the tab to see the rebuild status.

You can view the RAID rebuild rate and percent complete.

3. Click the link whenever you want to update the progress.

Return storage node to cluster

Return the repaired storage node to the cluster.

1. In the navigation window, right-click the storage node and select Add to Existing Management Group.

2. Select from the list the Group Name that the storage node used to belong to and click Add. The storage node 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 node is 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 node before you began the repair process, you may start a manager on the repaired storage node, 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.

- First, right-click on the virtual manager and select Stop Virtual Manager.
- Next, right-click on the virtual manager and select Delete Virtual Manager.
- Finally, right-click on the storage node and select Start Manager.

Add repaired node to cluster

1. After the initialization completes, right-click on the cluster and select Edit Cluster. The list of the storage nodes in the cluster should include the ghost IP address.

You now need to add the repaired storage node to the cluster in the spot held by the ghost IP address.

2. In the Edit Cluster window, first note the order of the storage nodes in the list.
3. Next, remove the ghost storage node from the cluster.
4. Return the repaired storage node to the cluster in the position of the ghost storage node.

Use the arrows to return the storage nodes in the list to their original order.

Example: If the list of the storage nodes in the cluster had storage node A, <IP address>, storage node C, and storage node B is the repaired storage node, then after re-arranging, the list of the storage nodes in the cluster should be storage node A, storage node B, storage node C and the storage node <IP address> will be in the management group.

Table 73 Replacing the ghost storage node with the repaired storage node

Storage Nodes in Cluster	
Before rearranging	<ul style="list-style-type: none"> • Storage node A • <IP Address> • Storage node C
After rearranging	<ul style="list-style-type: none"> • Storage node A • Storage node B • Storage node C

 **NOTE:**

If you do not arrange the storage nodes to match their original order, the data in the cluster is rebuilt across all the storage nodes instead of just the repaired storage node. 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 node goes through the rebuild, before you click the OK button in the Edit Cluster window, double-check that the order of the storage nodes in the cluster list matches the original order.

Rebuild volume data

After the storage node is successfully added back to the cluster, the adjacent storage nodes start rebuilding data on the repaired storage node.

1. Select the cluster and select the Disk Usage tab.
2. Verify that the disk usage on the repaired storage node 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 node.

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 node. 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 node

Remove the ghost storage node after the data is rebuilt.

The data is rebuilt on the storage node when two conditions are met.

- The repaired storage node's disk usage matches the usage of the other storage nodes in the cluster, and
- 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 node are successfully replaced, the data will be fully rebuilt on that storage node, 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 pre-paid packing slip in the replacement package. Put the RA number on the package as instructed by Customer Support.

22 iSCSI and the HP LeftHand Storage 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”](#) on page 335
- [“iSNS server”](#) on page 336
- [“iSCSI load balancing”](#) on page 336
- [“Authentication \(CHAP\)”](#) on page 337
- [“iSCSI and CHAP terminology”](#) on page 338
- [“About HP LeftHand DSM for MPIO”](#) on page 340

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”](#) on page 174.

Virtual IP addresses

A virtual IP (VIP) address is a highly available IP address which ensures that if a storage node in a cluster becomes unavailable, servers can still access a volume through the other storage nodes 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 SAN/iQ HP LeftHand DSM for MPIO.

When using a VIP, one storage node in the cluster hosts the VIP. All I/O goes through the VIP host. You can determine which storage node 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 nodes occupying the same cluster must be in the same subnet address range as the VIP.
- The VIP must be routable regardless of which storage node 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 node 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 nodes 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 node hosting the load balancing iSCSI session.

Configure iSCSI load balancing when setting up servers. See [Chapter 17](#) on page 289 for information about configuring server access for volumes.

Requirements

- Cluster configured with a virtual IP address. See “[Virtual IP addresses](#)” on page 335.
- A compliant iSCSI initiator.

Compliant iSCSI initiators

A compliant initiator supports iSCSI Login-Redirect AND has passed HP LeftHand Networks’ 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, [Figure 168](#).

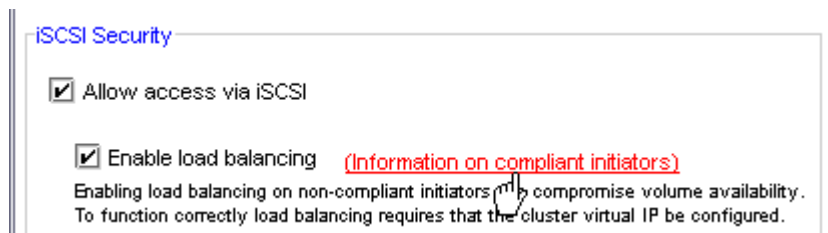


Figure 168 Finding compliant initiator information

The link opens the iSCSI initiator information window, [Figure 169](#). Scroll down for a list of compliant initiators.

If your initiator is not on the list, do not enable load balancing.

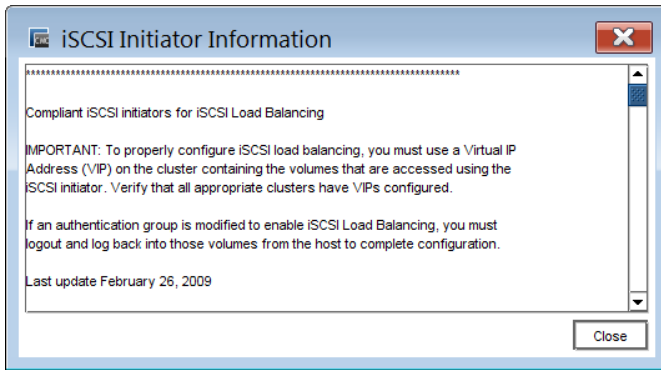


Figure 169 Viewing compliant iSCSI initiators

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](#)” on page 338.

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.

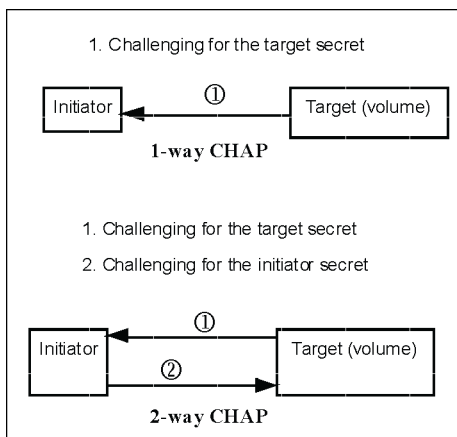


Figure 170 Differentiating types of CHAP

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 74](#) lists the requirements for configuring CHAP.

Requirements for configuring CHAP

Table 74 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 75 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 171 illustrates the configuration for a single host authentication with CHAP not required with Microsoft iSCSI.

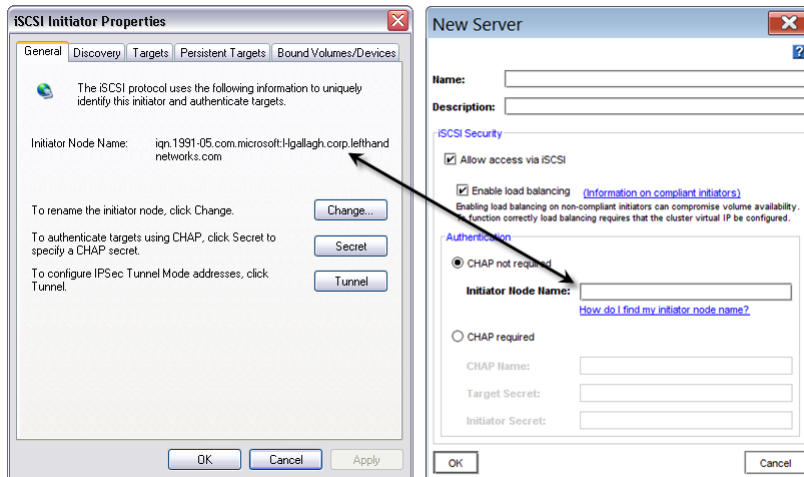


Figure 171 Viewing the MS iSCSI initiator to copy the initiator node name

Figure 172 illustrates the configuration for a single host authentication with 1-way CHAP required.

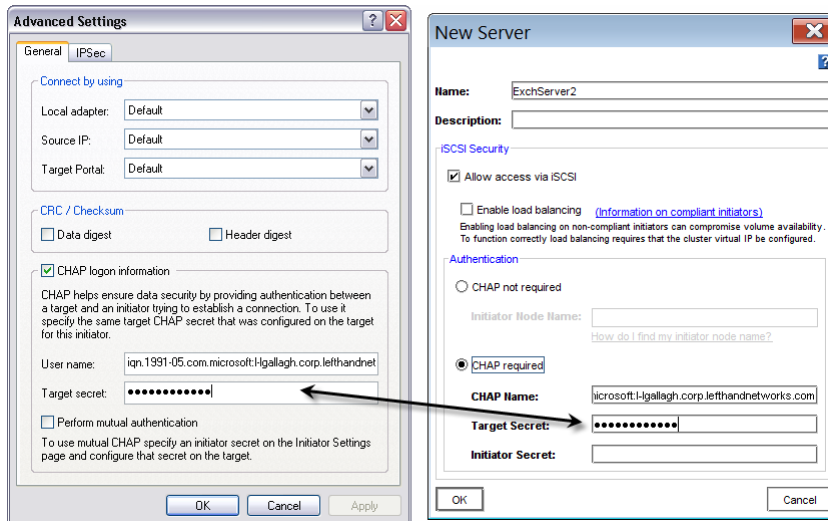


Figure 172 Configuring iSCSI (shown in the MS iSCSI initiator) for a single host with CHAP

Figure 173 illustrates the configuration for a single host authentication with 2-way CHAP required.

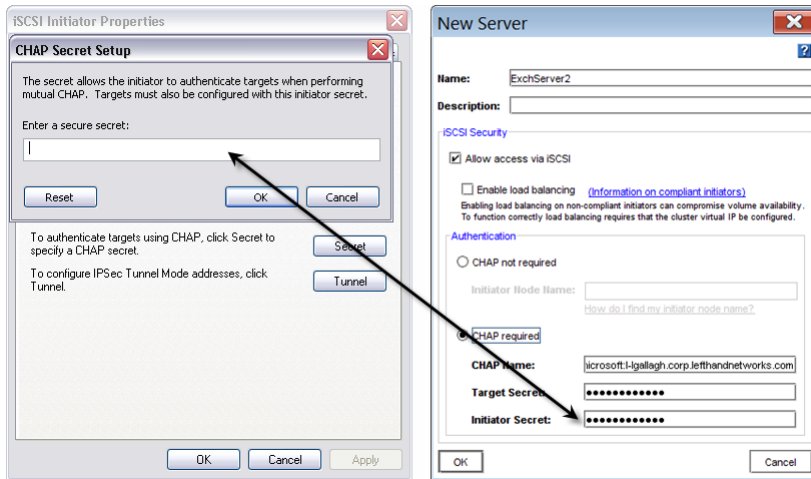


Figure 173 Adding an initiator secret for 2-way CHAP (shown in the MS iSCSI initiator)

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 LeftHand DSM for MPIO

If you are using the SAN/iQ HP LeftHand DSM for MPIO, you can use HP LeftHand DSM for MPIO to access volumes. For more information about HP LeftHand DSM for MPIO, refer to the HP LeftHand P4000 Windows Solution Pack User Manual.

You can see if you are using HP LeftHand 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 LeftHand DSM for MPIO, keep in mind the following:

- SAN/iQ HP LeftHand DSM for MPIO and the Microsoft MPIO must be installed on the server.
- With the above installed, servers automatically use HP LeftHand 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.

23 Using the Configuration Interface

The Configuration Interface is the command line interface that uses a direct connection with the storage node.

You may need to access the Configuration Interface if all network connections to the storage node are disabled. Use the Configuration Interface to perform the following tasks.

- Add storage node 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 node from a management group
- Reset the storage node 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 node 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 node 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 configura-
tion
# 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 mautobaud = Yes
pu backspace = DEL
pu hasdcd = No
pu rtscts = No
pu xonxoff = Yes
pu askndir = 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. Type `start` and press Enter at the log in prompt.
3. When the session is connected to the storage node, the Configuration Interface window opens.

Logging in to the Configuration Interface

Once you have established a connection to the storage node, log in to the Configuration Interface.

Table 76 Logging in depends on where the storage node is

If the storage node is in	From Configuration Interface entry window
Available Nodes pool	<ul style="list-style-type: none"> • 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. the Configuration Interface Login window opens. 2. Type the user name and password of the administrative user created for the management group. 3. Tab to Login and press Enter. The Configuration Interface main menu opens.

**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 node comes with two Ethernet interfaces. [Table 77](#) lists where interfaces are labeled and the label name.

Table 77 Identifying ethernet interfaces on the storage node

Ethernet Interfaces	
Where labeled	What the label says
TCP/IP Network configuration category in the CMC <ul style="list-style-type: none"> • TCP/IP tab • TCP Status tab 	Name: eth0, eth1 Motherboard:Port0, Motherboard:Port1 G4-Motherboard:Port1, G4-Motherboard:Port2 Motherboard:Port1, Motherboard:Port2
Configuration Interface	Intel Gigabit Ethernet or Broadcom Gigabit Ethernet
Label on the back of the storage node	Eth0, Eth1 or represented by a graphical symbol similar to the symbols below:

Once you have established a connection to the storage node 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, the 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.
A window opens listing the assigned IP address.
6. Open the CMC and locate the storage node using the Find function.

Deleting a NIC bond

You can delete 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](#)” on page 92.

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 a 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.
The Available Network Devices window opens. The logical bond is the only interface listed in the window.
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 node 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](#)” on page 108.

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”](#) on page 109.

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 type 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 node from a management group

Removing a storage node from a management group deletes all data from the storage node, clears all information about the management group, and reboots the storage node.

△ CAUTION:

Removing a storage node from a management group deletes all data on the storage node.

1. On the Configuration Interface main menu, tab to Config Management and press Enter.
2. Tab to Remove from management group and press Enter.

A window opens, warning you that removing the storage node from the management group will delete all data on the storage node and reboot it.

3. Tab to Ok and press Enter
4. On the Configuration Management window, tab to Done and press Enter.

Resetting the storage node to factory defaults

Resetting the storage node to factory defaults deletes all data and erases the configuration of the storage node, including administrative users and network settings.

△ CAUTION:

Resetting the storage node to factory defaults deletes all data on the storage node.

1. On the Configuration Interface main menu, tab to Config Management and press Enter.

2. Tab to Reset to factory defaults and press Enter.

A window opens, warning you that resetting the storage node configuration will delete all data on the storage node and reboot the storage node.

3. Tab to Ok and press Enter.

4. On the Configuration Management window, tab to Done and press Enter.

24 Glossary

Terms Used

The following glossary provides definitions of terms used in the SAN/iQ software and the HP LeftHand Storage Solution.

Table 78 Glossary

Term	Definition
Active Monitoring	Active monitoring tracks the health of the storage node using notifications such as emails, alerts in the CMC, and SNMP traps.
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.
Add-on application	An additional feature purchased separately from the SAN/iQ software.
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 nodes on the subnet the CMC is connected to. Any storage nodes 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 node boots up. Also known as disk-on-modules or DOMs.
HP LeftHand Centralized Management Console	Management interface for the SAN/iQ software.
CHAP	Challenge-Handshake Authentication Protocol (CHAP) is a standard authentication protocol.
Clone point	The snapshot that has two or more volumes associated with it. A clone point is created when a SmartClone volume is created from a snapshot or from snapshot temporary space.

Term	Definition
CLI	Command-line interface for the SAN/iQ software.
Cluster	A cluster is a grouping of storage nodes that create the storage pool from which you create volumes.
CMC	Centralized Management Console. See HP LeftHand Centralized Management Console.
Communication Mode	The unicast communication among storage nodes 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 nodes, and iSCSI sessions in the HP LeftHand Storage Solution. It provides an overview of the storage network broken out by management groups.
Date and Time	Set the date and time on the storage node if not using an external time service, such as NTP.
Diagnostics	Diagnostics check the health of the storage node hardware.
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
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 NSM	When using Repair Storage Node, a “ghost” storage node acts as a placeholder in the cluster, keeping the cluster intact, while you repair or replace the storage node.
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 node, its drives, and configuration.
Hostname	The hostname on a storage node is the user-definable name that displays below the storage node icon in the network window. It is also visible when the users browse the network.
ID LED	LED lights on the physical storage node so that you can find that node in a rack [NSM 260 only].
iSCSI Load Balancing	Improves iSCSI performance and scalability by distributing iSCSI sessions for different volumes evenly across storage nodes in a cluster.

Term	Definition
License keys	A license key registers a storage node for add-on applications. Each storage node requires its own license key.
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 node are stored both locally on the storage node and are also written to a remote log server.
Management group	A collection of one or more storage nodes which serves as the container within which you cluster storage nodes and create volumes for storage.
Managers	Manager software runs on storage nodes within a management group. You start managers on designated storage nodes to govern the activity of all of the storage nodes in the group.
MIB	The Management Information Base provides SNMP read-only access to the storage node.
Monitored variables	Variables that report health status of the storage node. These variables can be monitored using alerts, emails, and SNMP traps.
Network window	Graphically depicts the status of each storage node. Storage Nodes on the network are either available or part of a management group.
NTP	Network Time Protocol
Parity	In RAID5, redundant information is stored as parity distributed across the disks. Parity allows the storage node to use more disk capacity for data storage.
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 associated with the cluster.
point-in-time consistent 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.
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	Type of RAID configuration. <ul style="list-style-type: none"> • RAID0 - data striped across disk set • RAID1 - data mirrored from one disk onto a second disk • RAID10 - mirrored sets of RAID1 disks • RAID5 - data blocks are distributed across all disks in a RAID set. Redundant information is stored as parity distributed across the disks. • RAID50 - mirrored sets of RAID5 disks.

Term	Definition
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 node</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 node. The storage node is offline and flashes red in the network window.
Register	Register individual storage nodes to use add-on applications. Registration requires sending in the storage node serial numbers to purchase the license keys, which are then applied to the storage node.
Repair NSM	Creates a placeholder in the cluster, in the form of a “ghost” storage node, that keeps the cluster intact while you remove the storage node to replace a disk or replace the storage node itself, and return it to the cluster.
Replication Level	Designate how many copies of data you want to keep in the cluster.
Replication Priority	Designates whether data availability or redundancy is more important in your configuration.
Restripe	Striped data is stored across all disks in the cluster. You might change the configuration of a volume, for example, change replication level, add a storage node, or remove a storage node. 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 node goes down, and writes continue to a second storage node, and the original store comes back up, the original storage node needs to recoup the exact data captured by the second storage node.
Rolling Back	Replaces the original volume with a read/write copy of a selected snapshot. New for release 8.0: The new volume retains the same name.
SAN/iQ Interface	When you initially set up a storage node using the Configuration Interface, the first interface that you configure becomes the interface used for the SAN/iQ software communication.
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.
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	A fixed version of a volume for use with backup and other applications.

Term	Definition
snapshot set	Application-managed snapshots created for a volume set.
SNMP Traps	Use traps to have an SNMP tool send alerts when a monitoring threshold is reached.
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.
Target Secret	Target secret is used in both 1-way and 2-way CHAP when the target (volume) challenges the iSCSI initiator.
Thin Provisioning	Thin provisioning reserves less space on the SAN than is presented to application servers.
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.
Time Zone	The time zone for the physical location of the storage node.
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.
Virtual IP Address	A highly available address that ensures that if a storage node in a cluster becomes unavailable, servers can still access the volume through the other storage nodes 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 node until it is needed to regain quorum.
Volume	A logical entity that is made up of storage on one or more storage nodes. 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 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 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 Size	The size of the virtual device communicated to the operating system and the applications.
VSS	Volume Shadow Copy Service
VSS Provider	HP LeftHand P4000 VSS Provider is the hardware provider that supports the Volume Shadow Copy Service on the HP LeftHand Storage Solution.
Writable Space	See Temporary Space.

Index

Symbols

- 1000BASE T interface, 90
- 30-day evaluation for add-on applications, 317
- 3650 See IBM x3650, 55

A

- access control
 - SNMP, 130
- Access Rights See Permission Levels, 293
- Access Volume wizard
 - servers, 39
- accessing
 - volumes from servers, 289
- activating dedicated boot devices, 54
- active interface
 - in active-passive bond, 95
 - in adaptive load balancing bond, 99
 - in link aggregation dynamic mode bond, 97
- active monitoring, 135
- Active-Passive bond, 94
 - active interface, 95
 - during failover, 95
 - example configurations, 96
 - requirements, 95
- Adaptive Load Balancing bond
 - active interface, 99
 - during failover, 99
 - example configurations, 100
 - preferred interface, 99
 - requirements, 98
- add-on applications
 - evaluating, 317
 - overview, 317
 - registering for, 320
- Adding
 - servers to management groups, 290
 - statistics, 310
- adding
 - a remote log, 168
 - administrative groups, 125
 - administrative users, 123
 - clusters, 209
 - adding storage to, 212
 - DNS domain name, 112
 - DNS servers, 112
 - domain names to DNS suffixes, 113
 - iSNS server, 210
 - management groups, 177
 - requirements for, 178
 - managers to management group, 181
 - monitored variables, 136
 - routes, 114
 - snapshot schedules, 255
 - snapshots, 248
 - requirements, 246
 - SNMP clients, 130
 - storage for the first time, 37
 - storage nodes to existing cluster, 211
 - storage nodes to management group, 180
 - users to a group, 127
 - virtual manager, 204
 - volumes, 239
 - requirements for, 238
- administrative groups, 125
 - adding, 125
 - adding users, 127
 - changing, 126
 - deleting, 127
 - permission levels, 126
 - permissions descriptions, 126, 127
 - removing users, 127
- administrative security, 171
- administrative users, 123
 - adding, 123
 - deleting, 124
- agents, disabling SNMP, 132
- alert notification
 - via email, 141
 - via SNMP, 141
 - via the CMC, 141

- alerts
 - active monitoring, [135](#)
 - editing variables in alerts, [137](#)
 - selecting alerts to monitor, [136](#)
 - tab window, [136](#)
 - viewing and saving, [143](#)
 - window for viewing, [30](#)
- anager IP addresses
 - updating, [117](#)
- application-managed snapshots
 - converting temporary space from, [253](#)
 - creating, [248](#)
 - creating for volume sets, [249](#)
 - creating SmartClone volumes from, [260](#)
 - defined, [245](#)
 - deleting, [261](#)
 - making available, [250](#), [251](#), [252](#)
 - requirements for, [248](#)
 - rolling back from, [259](#)
- Assigning servers to volumes and snapshots, [292](#), [293](#), [294](#)
- Authentication Groups
 - and volume lists, [289](#)
- auto discover, [29](#)
- auto performance protection, [214](#)
 - storage server inoperable, [214](#)
 - storage server overloaded, [214](#)
 - volume availability and, [214](#)
- availability of volumes and snapshots, [51](#), [214](#)
- Availability tab, [51](#)
- available node pool, [171](#)
- available nodes, [32](#)
- available storage nodes
 - See Available Nodes

B

- backing out
 - of Remote Copy evaluation, [318](#)
 - of scripting evaluation, [320](#)
- backing up
 - management group configuration description, [184](#)
 - management group configuration, binary, [183](#)
 - Management Group with Remote Copy Relationship, [184](#)
 - storage node configuration file, [46](#)
- backup and restore
 - storage node configuration files, [45](#)
- bandwidth, changing local settings, [183](#)
- benefits of virtual manager, [201](#)

- best practice
 - recommended numbers for management
 - group storage items,
 - configuring cluster for disaster recovery, [202](#)
 - configuring volumes with critical data, [225](#)
 - frame size, [110](#)
 - link aggregation dynamic mode, [93](#)
 - NIC bonds, [100](#)
 - provisioning storage, [222](#)
 - setting replication levels and redundancy modes, [225](#)
 - setting volume size, [222](#)
 - speed and duplex settings, [109](#)
 - using snapshots
 - as source volumes for data mining, [247](#)
 - for data preservation, [247](#)
 - protection against data deletion, [247](#)
- block device, iSCSI as, [233](#)
- block storage system, [233](#)
- boot devices
 - activating dedicated, [54](#)
 - checking status of dedicated, [52](#)
 - configuring dedicated, [53](#)
 - deactivating dedicated, [53](#)
 - dedicated, [51](#)
 - removing dedicated, [53](#)
 - replacing dedicated, [53](#)
 - starting and stopping dedicated, [52](#)
 - status of dedicated, [52](#)
- Boot Devices tab, [52](#)
- BOOTP, [91](#)

C

- Capacity
 - RAID50, [59](#)
- capacity
 - clusters, [209](#)
 - disk capacity and volume size, [233](#)
 - of the SAN, [221](#)
 - planning volume size, [221](#)
 - planning, full provisioning, [222](#)
 - planning, thin provisioning, [222](#)
 - RAID0, [57](#)
 - RAID10, [57](#)
 - RAID5, [58](#)
 - RAID6, [59](#)
 - storage nodes, [209](#)
- capacity management
 - and scheduled snapshots, [254](#)
 - snapshot thresholds, [226](#)

- Centralized Management Console
 - overview,
 - alerts window, [34](#)
 - features of ,
 - Getting Started Launch Pad, [32](#)
 - icons used in ,
 - menu bar, [30](#)
 - navigation window, [31](#)
 - tab window, [33](#)
- Challenge Handshake Authentication Protocol
 - See CHAP
- changing
 - administrative group description, [126](#)
 - administrative group permissions, [126](#)
 - changing RAID erases data, [70](#)
 - cluster configuration, [211](#)
 - clusters for volumes, [242](#)
 - data availability, [242](#)
 - data redundancy, [242](#)
 - host names, [44](#)
 - IP address of storage node, [91](#)
 - local bandwidth, [183](#)
 - maintenance mode to normal, [186](#)
 - management groups, [182](#)
 - order of NTP server access, [121](#)
 - replication levels, [242](#)
 - replication priority, [242](#)
 - snapshots, [250](#)
 - thresholds in a snapshot, [250](#)
 - user password, [124](#)
 - volume descriptions, [241](#)
 - volume size, [242](#)
- CHAP
 - 1-way, [337](#)
 - 2-way, [337](#)
 - editing, [291](#)
 - iSCSI, [337](#)
 - requirements for configuring, [291](#), [338](#)
 - terminology in different initiators, [338](#)
 - using, [337](#)
 - volumes and, [337](#)
- characteristics of SmartClone volumes, [267](#)
- checklist for disk replacement, [82](#)
- choosing a RAID configuration, [57](#)
- clearing
 - items in navigation window, [40](#)
 - statistics sample data, [312](#)
- client access to volumes using Access Volume wizard, [39](#)
- clients, adding SNMP, [130](#)
- clone
 - See SmartClone volumes
- clone a volume, [265](#)
- clone point
 - and shared snapshots, [274](#)
 - deleting, [285](#)
 - utilization of, [282](#)
- clone point ,
- clustering managers, [171](#)
- Clusters
 - comparing the load of two, [301](#), [308](#)
- clusters
 - adding, [209](#)
 - adding storage node, [211](#)
 - capacity, [211](#)
 - changing volumes, [242](#)
 - data replication levels, [223](#)
 - defined, [32](#)
 - deleting, [219](#)
 - editing, [211](#)
 - overview, [209](#)
 - prerequisites for, [209](#)
 - removing storage nodes from, [212](#)
 - repairing storage node in, [216](#)
 - troubleshooting, [214](#)
- CMC
 - see Centralized Management Console,
- communication interface for SAN/iQ
- communication, [116](#)
- community strings
 - reserved on the DL 320s, [130](#)
 - reserved on the DL 380, [130](#)
- Compliant iSCSI Initiators, [290](#), [336](#)
- compliant iSCSI initiators, [336](#)
- configuration categories
 - storage node, defined, [43](#)
 - storage nodes, [43](#)
- configuration file
 - backing up management group configuration, [184](#)
 - backing up storage node configuration, [46](#)
 - restoring, [46](#)
- Configuration Interface
 - :deleting NIC bond in, [344](#)
 - configuring frame size in, [344](#)
 - configuring network connection in, [343](#)
 - configuring TCP speed and duplex in, [344](#)
 - connecting to, [341](#)
 - creating administrative users in, [343](#)
 - resetting DSM configuration in, [345](#)
 - resetting NSM to factory defaults in, [345](#)
- Configuration Summary
 - overview, [174](#)
 - reading for management group, [175](#)
- configuration summary
 - configuration guidance, [174](#)
 - management group summary roll-up, [174](#)

- configurations
 - RAID, [57](#)
- Configuring
 - iSCSI single host, [339](#)
- configuring
 - dedicated boot devices, [53](#)
 - disabled network interface, [107](#)
 - Failover Managers, [190](#)
 - frame size in Configuration Interface, [344](#)
 - IP address manually, [91](#)
 - multiple storage nodes, [40](#)
 - network connection in Configuration Interface, [343](#)
 - network interface bonds, [100](#)
 - network interfaces, [91](#)
 - NIC, speed and duplex, [108](#)
 - RAID, [56](#)
 - storage nodes, [37](#)
 - TCP speed and duplex in Configuration Interface, [344](#)
 - virtual IP address, [210](#)
 - virtual manager, [204](#)
- connecting to the Configuration Interface, [341](#)
- converting temporary space
 - from application-managed snapshots, [253](#)
- copying storage node configurations, [40](#)
- copying, volumes from snapshots, [250](#)
- creating
 - See adding
 - administrative users in Configuration Interface, [343](#)
 - new administrative group, [127](#)
 - NIC bond, [101](#)
 - SmartClone volumes, [276](#)
 - volumes using the wizard, [38](#)
- creating storage, [221](#)
- CSV file, exporting performance statistics to, [315](#)
- duplex, configuring, [108](#)
- customer information, [324](#)
- customer support
 - registering Remote Copy, [320](#)

D

- data
 - and deleting volumes,
 - clearing statistics sample, [312](#)
 - preserving with snapshots, [247](#)
- data availability
 - changing, [242](#)
 - requirements for setting replication priority, [239](#)
- data mining using SmartClone volumes, [265](#)
- data redundancy
 - and RAID status, [71](#)
 - changing, [242](#)
 - requirements for setting replication priority, [239](#)
- data replication, [223](#)
 - levels allowed in clusters, [223](#)
 - planning, [223](#)
 - requirements for setting, [238](#)
- data transfer and RAID status, [71](#)
- data transmission, [111](#)
- date
 - setting with NTP, [120](#)
 - setting without NTP, [121](#)
- date and time for scheduled snapshot, [255](#)
- decreasing volume size, [242](#)
- defaults
 - restoring for the Performance Monitor, [313](#)
- definition
 - clusters, [32](#)
 - management groups, [32](#)
 - remote copies, [32](#)
 - servers, [32](#)
 - sites, [32](#)
 - SmartClone volumes, [263](#)
 - snapshots, [32](#)
 - volumes, [32](#)
- definition of
 - RAID configurations, [57](#)
- degraded RAID, [71](#)
- Deleting
 - servers, [292](#)
- deleting
 - administrative groups, [127](#)
 - administrative users, [124](#)
 - an administrative group, [127](#)
 - clone point, [285](#)
 - clusters, [219](#)
 - DNS servers, [114](#)
 - management groups, [187](#)
 - multiple SmartClone volumes, [285](#)
 - network interface bonds, [104](#)
 - NIC bond in Configuration Interface, [344](#)
 - NTP server, [121](#)
 - routing information, [115](#)
 - SmartClone volumes, [284](#)
 - snapshot schedules, [256](#)
 - snapshots, [227](#), [261](#)
 - volumes, [243](#)
 - prerequisites for, [258](#), [261](#)

- Dell 2950
 - capacity
 - RAID5, [58](#)
 - disk arrangement in,
 - disk setup, [77](#)
 - disk status, [77](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - RAID10 initial setup, [61](#)
 - RAID5 initial setup, [63](#)
- descriptions
 - changing for clusters, [211](#)
 - changing for volumes, [241](#)
- Details tab
 - storage nodes, [49](#)
- Details, viewing for statistics, [312](#)
- DHCP
 - using, [91](#)
 - warnings when using, [92](#)
- diagnostics
 - hardware, [144](#)
 - list of diagnostic tests, [145](#)
 - viewing reports, [145](#)
- disabled network interface, configuring, [107](#)
- disabling
 - network interfaces, [106](#)
 - SNMP agent, [132](#)
 - SNMP traps, [134](#)
- Disassociating Management Groups
 - See Remote Copy Users Guide
- disaster recovery
 - best practice, [202](#)
 - starting virtual manager, [205](#)
 - using a virtual manager, [200](#)
- disk
 - arrangement in platforms, [74](#)
 - disk setup report, [72](#)
 - managing, [72](#)
 - managing in storage node, [72](#)
 - powering off through the CMC, [83](#), [84](#)
 - powering on through the CMC, [83](#), [85](#)
 - replacement, [83](#), [85](#), [86](#)
 - replacement checklist, [82](#)
 - replacing in RAID1/10 and 5/50, [83](#)
 - replacing in replicated cluster, [216](#)
 - replacing in storage node, [80](#)
 - using Repair Storage Node when replacing, [80](#)
 - VSA, recreating, [77](#)
- disk drive
 - see disk, [73](#)
- disk report, [73](#)
- disk setup
 - Dell 2950, [77](#)
 - DL 380, [75](#)
 - DL320s, [76](#)
 - HP LeftHand P4300, [80](#)
 - HP LeftHand P4500, [79](#)
 - NSM 2060, [77](#)
 - NSM 4150, [78](#)
 - report, [73](#)
 - tab, [73](#)
- disk space usage, [233](#)
- disk status
 - Dell 2950, [77](#)
 - DL320s, [76](#)
 - DL380, [75](#)
 - HP LeftHand P4500, [79](#)
 - IBM x3650, [77](#)
 - NSM 160, [74](#)
 - NSM 2060, [77](#)
 - NSM 260, [75](#)
 - NSM 4150, [78](#)
 - VSA, [77](#)
- DL 320s
 - reserved community strings, [130](#)
- DL 380
 - reserved community strings, [130](#)
- DL320s
 - capacity
 - RAID5, [58](#)
 - capacity in RAID6, [59](#)
 - disk setup, [76](#)
 - disk status, [76](#)
 - drive failure, [59](#)
 - hot swapping, [59](#)
 - in RAID10, [61](#)
 - parity in RAID6, [59](#)
 - RAID rebuild rate, [69](#)
 - RAID5 set, [63](#)
 - RAID6, [66](#)
- DL380
 - capacity
 - RAID5, [58](#)
 - disk arrangement in,
 - disk setup, [75](#)
 - disk status, [75](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - RAID0, [60](#)
 - RAID10, [61](#)
- DNS
 - adding domain name, [112](#)
 - and DHCP, [112](#)
 - and static IP addresses, [112](#)

- DNS server
 - adding, [112](#)
 - and static IP addresses, [112](#)
 - editing IP or domain name, [113](#)
 - removing, [114](#)
 - using, [112](#)
- documentation
 - HP website, [27](#)
 - providing feedback, [28](#)
- Domain Name Server
 - See DNS Server
- domain names
 - adding to DNS suffixes, [113](#)
 - editing in DNS suffixes list, [113](#)
 - removing from DNS suffixes list, [114](#)
- downloading variable log file, [143](#), [144](#)
- DSM
 - when using two NICs, [295](#)
- DSM for MPIO, [335](#)
 - how to see if using, [340](#)
 - tips for using to access volumes from servers, [340](#)
- Dynamic Host Configuration Protocol
 - See DHCP

E

- Editing
 - servers, [291](#)
- editing
 - clusters, [211](#)
 - DNS server domain names, [113](#)
 - DNS server IP addresses, [113](#)
 - domain name in DNS suffixes list, [113](#)
 - frame size, [110](#)
 - group name, [126](#)
 - management groups, [182](#)
 - monitored variables, [137](#)
 - network interface
 - frame size, [109](#)
 - speed and duplex, [108](#)
 - NTP server, [120](#)
 - routes, [115](#)
 - SmartClone volumes, [284](#)
 - snapshot schedules, [255](#)
 - snapshots, [250](#)
 - SNMP trap recipient, [133](#)
 - volumes, [240](#)
- enabling
 - NIC flow control, [111](#)
 - SNMP Traps, [133](#)
- establishing network interfaces, [90](#)
- eth0 and eth1, [90](#)
- ethernet interfaces, [90](#)

- evaluating
 - add-on applications, [317](#)
 - Remote Copy, [318](#)
 - backing out of, [318](#)
 - scripting, [319](#)
 - backing out of, [320](#)
- example scenarios for using SmartClone volumes, [264](#)
- Exporting
 - performance data, [315](#)
 - performance statistics to a CSV file, [315](#)

F

- Failover Manager, [173](#)
 - and Multi-Site SAN, [173](#)
 - configuring, [190](#)
 - overview, [189](#)
 - requirement for using, [190](#)
 - requirements for, [189](#)
 - troubleshooting, [194](#)
 - using in Multi-Site SAN, [189](#)
- fault tolerance, [335](#)
 - network interface bonding, [92](#)
 - quorum and managers, [172](#)
 - replication level for volumes, [223](#)
 - replication priority for volumes, [225](#)
 - stopping managers, [182](#)
- Faults, isolating, [299](#)
- feature registration, [180](#)
- Feature Registration tab, [320](#), [321](#)
- features
 - of Centralized Management Console, [29](#)
- file systems, [233](#)
 - mounting on volumes, [237](#)
- find Failover Manager on network, [195](#)
- finding SNMP MIB, [132](#)
- finding storage nodes, [37](#)
 - Auto Discover, [29](#)
 - on the network, [39](#)
- first storage node, [178](#)
- flow control, [111](#)
- Formatting volumes for use, [295](#)
- frame size, NIC, [109](#)
- frame size, VSA, [89](#)
- frames, editing size, [110](#)
- full permissions, [126](#)

G

- gateway session for VIP with load balancing, [336](#)
- Getting Started Launch Pad, [37](#)
- ghost storage node, [216](#)
 - removing after data is rebuilt, [333](#)

- Gigabit Ethernet, [90](#)
 - See also GBe
- Glossary
 - for SAN/iQ software and LeftHand SAN, [347](#)
- glossary
 - SmartClone volumes, [264](#)
- Graphical Legend
 - Hardware tab, [33](#)
- group name
 - editing, [126](#)
- groups
 - administrative default groups, [125](#)
 - deleting administrative, [127](#)
- groups, administrative, [125](#)

H

- hardware diagnostics, [144](#)
 - list of diagnostic tests, [145](#)
 - tab window, [144](#)
- hardware information
 - log files, [167](#)
- hardware information report, [144](#), [150](#)
 - details, [152](#)
 - expanded details, [150](#)
 - generating, [150](#)
 - saving to a file, [151](#)
 - updating, [150](#)
- Hardware tab in graphical legend, [33](#)
- help
 - obtaining, [27](#)
- Highlighting lines, [315](#)
- host names
 - access SNMP by, [130](#)
 - changing, [44](#)
 - resolution, [44](#)
- host storage node for virtual IP address, [335](#)
- hot spares
 - in RAID5, [58](#)
- hot swap, [59](#)
 - RAID degraded ,
 - safe to remove status, [81](#)
- HP
 - technical support, [27](#)
- HP LeftHand P4300
 - capacity
 - RAID5, [58](#)
 - capacity:RAID6, [59](#)
 - disk setup, [80](#)
 - parity in RAID6, [59](#)
 - RAID levels and default configuration, [55](#)

- HP LeftHand P4500
 - capacity
 - RAID5, [58](#)
 - capacity in RAID6, [59](#)
 - disk setup, [79](#)
 - disk status, [79](#)
 - drive failure, [59](#)
 - hot swapping, [59](#)
 - in RAID10, [61](#)
 - parity in RAID6, [59](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - RAID5 set, [63](#)
 - RAID6, [66](#)
- HP System Insight Manager (HP SIM), logging into, [130](#)

I

- I/O performance, [214](#)
- IBM x3650
 - capacity
 - RAID5, [58](#)
 - disk arrangement in,
 - RAID levels and default configuration, [55](#)
- icons
 - licensing, [318](#)
 - used in Centralized Management Console, [33](#)
- identifying network interfaces, [90](#)
- increasing volume size, [242](#)
- Insight Manager, [129](#)
- installing Failover Manager, [191](#)
- installing SNMP MIB, [132](#)
- interface
 - administrative users in, [343](#)
 - configuration, [341](#)
 - configuring network connection in, [343](#)
 - connecting to, [341](#)
 - deleting NIC bond in, [344](#)
 - resetting NSM to factory defaults, [345](#)
- IP addresses
 - accessing SNMP by, [130](#)
 - changing, iSNS server ,
 - configuring for storage node, [91](#)
 - NTP server, [120](#)
 - removing, iSNS server, [213](#)
 - using DHCP/BOOTP, [91](#)

- iSCSI
 - and CHAP, [337](#)
 - and fault tolerance, [335](#)
 - and iSNS servers, [336](#)
 - and virtual IP address, [335](#)
 - as block device, [233](#)
 - authentication, [337](#)
 - changing or removing virtual IP, [212](#)
 - CHAP, [337](#)
 - clusters and VIP, [335](#)
 - configuring CHAP, [291](#), [338](#)
 - load balancing, [290](#), [336](#)
 - load balancing and compliant initiators, [290](#), [336](#)
 - logging on to volumes, [295](#)
 - performance, [336](#)
 - setting up volumes as persistent targets, [295](#)
 - single host configuration, [339](#)
 - terminology in different initiators, [338](#)
 - volumes and, [337](#)
- iSCSI initiators
 - configuring virtual IP addresses for, [210](#)
- iSNS Server
 - and iSCSI targets, [336](#)
- iSNS server
 - adding, [210](#)
- iSNS server
 - changing or removing IP address,

L

- layout of disks in platforms, [74](#)
- license information, [324](#)
- license keys, [320](#)
- licensing icons, [318](#)
- Lines
 - changing the color of in the Performance Monitor, [314](#)
 - changing the style of in the Performance Monitor, [314](#)
 - displaying or hiding in the Performance Monitor, [314](#)
 - highlighting, [315](#)
- Link Aggregation Dynamic Mode bond, [97](#)
 - active interface, [97](#)
 - during failover, [98](#)
 - example configurations, [98](#)
 - preferred interface, [97](#)
 - requirements, [97](#)
- list of diagnostic tests, [145](#)
- list of monitored variables, [135](#)
- Load Balancing
 - compliant iSCSI initiators, [290](#)
 - editing, [291](#)
 - iSCSI, [290](#)

- load balancing
 - compliant iSCSI initiators, [336](#)
 - gateway session when using, [336](#)
 - iSCSI, [336](#)
- local bandwidth, setting, [183](#)
- locate NSM 260 in rack, [45](#)
- locating a storage node in a rack, [45](#)
- log files
 - backing up management group configuration file, [184](#)
 - backing up storage node configuration file, [46](#)
 - downloading variable, [143](#), [144](#)
 - hardware information, [167](#)
 - saving for technical support, [167](#)
- log in
 - to a storage node in a management group, [44](#)
 - to HP SIM, [130](#)
 - to management group, [180](#)
 - to storage nodes in Available Nodes pool, [31](#)
 - to System Management Homepage, [130](#)
- log out
 - of management group, [181](#)
- Logging on to volumes in iSCSI, [295](#)

M

- maintenance mode
 - changing to normal mode, [186](#)
 - management group in, [186](#)
- management group
 - registering, [180](#)
 - starting up, [185](#)
- management group time
 - refreshing, [119](#)
- Management Groups
 - adding servers to, [290](#)

- management groups
 - adding, 177
 - requirements for, 178
 - adding storage nodes, 178, 180
 - backing up configuration, 183
 - best practice recommendations, 175
 - configuration guidance, 174
 - configuration summary roll-up, 174
 - creating, 178
 - defined, 32
 - deleting, 187
 - editing, 182
 - function, 171
 - functions of managers, 172
 - logging in, 180
 - logging out, 181
 - maintenance mode, 186
 - normal mode, 186
 - overview, 171
 - reading configuration summary, 175
 - removing storage nodes, 187
 - prerequisites, 187
 - restoring, 184
 - setting local bandwidth, 183
 - shut down procedure, 185
 - shutting down, 184
 - starting managers, 181
 - starting up, 185
 - using virtual manager
 - configuration for, 200
 - disaster recovery, 200
- Management Information Base
 - See MIB
- managers
 - configuring Failover Manager, 190
 - Failover, 173
 - functions of, 172
 - implications of stopping, 182
 - overview, 171
 - quorum and fault tolerance, 172
 - starting, 181
 - stopping, 181
 - virtual, 200
- managing disks, 72
- Map View, 279
 - changing the view, 280
 - for SmartClone volumes, 280
 - toolbar, 280
- menu bar, 30

- MIB
 - exceptions, 325
 - for SNMP, 132
 - installing, 132
 - locating, 132
 - SNMP, 325
 - supported MIBs, 325
 - versions, 132
- migrating RAID, 69
- mixed RAID, 69
- monitored variables
 - adding, 136
 - editing, 137
 - list of, 135
 - removing, 138
 - viewing summary of, 141
- Monitoring
 - performance, 297
- Monitoring interval in the Performance Monitor, 309
- monitoring RAID status
 - status, 71
- Motherboard Port1 and Motherboard Port2, 90
- mounting snapshots, 250
- Multi-Site SAN
 - and Failover Manager, 173
 - using Failover Manager in, 189
- Multi-Site SAN sites, 32

N

- naming SmartClone volumes, 266, 267
- navigation window, 30
 - clearing items in, 40
- network
 - finding storage nodes on, 29, 39
 - managing settings, 89
 - overview, 89
 - settings for Failover Manager, troubleshooting, 195
- Network Interface Bonds
 - deleting, 344
 - determining of use would improve performance, 301

- network interface bonds, [92](#)
 - active-passive, [94](#)
 - adaptive load balancing, [98](#)
 - best practices, [100](#)
 - communication after deleting, [105](#)
 - configuring, [100](#)
 - creating, [101](#)
 - deleting, [104](#)
 - link aggregation dynamic mode, [97](#)
 - physical and logical interface, [94](#)
 - requirements, [93](#)
 - adaptive load balancing, [98](#)
 - setting flow control, and, [111](#)
 - settings, [93](#)
 - status of, [103](#)
 - verifying, [102](#)
 - VSA, [89](#)
- network interfaces, [97](#)
 - attaching Ethernet cables, [90](#)
 - bonding, [92](#)
 - configuring, [91](#), [107](#)
 - disabling or disconnecting, [106](#)
 - establishing, [90](#)
 - identifying, [90](#)
 - speed and duplex settings, [108](#)
 - used for SAN/iQ communication, [116](#)
 - VSA, [89](#)
- Network Time Protocol
 - See NTP
- network window
 - see navigation window, [31](#)
- NIC
 - See Network Interfaces, [90](#)
- NIC flow control, [111](#)
 - enabling, [111](#)
 - requirements, [111](#)
 - VSA, [89](#)
- nodes
 - finding on network, [29](#), [39](#)
- normal, [46](#)
- normal RAID
 - status, [71](#)
- not preferred NTP server, [120](#)
- NSM 160
 - capacity
 - RAID5, [58](#)
 - in RAID0, [60](#)
 - in RAID10, [61](#)
 - in RAID5, [63](#)
 - mirroring, [61](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - reconfiguring RAID, [71](#)
- NSM 2060
 - capacity
 - RAID5, [58](#)
 - disk arrangement in,
 - disk setup, [77](#)
 - disk status, [77](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - RAID10 initial setup, [61](#)
 - RAID5 initial setup, [63](#)
- NSM 260
 - capacity
 - RAID5, [58](#)
 - devices RAID10, [61](#)
 - disk status, [75](#)
 - locating in rack, [45](#)
 - RAID levels and default configuration, [55](#)
 - RAID rebuild rate, [69](#)
 - RAID0 devices, [60](#)
 - RAID50, [65](#)
- NSM 4150
 - powering off the system controller and disk enclosure, correct order,
 - powering on the system controller and disk enclosure, correct order,
 - disk arrangement in,
 - disk setup, [78](#)
 - disk status, [78](#)
 - RAID levels and default configuration, [55](#)
 - RAID10 initial setup, [61](#)
 - RAID50 capacity, [59](#)
 - RAID50 initial setup, [63](#)
 - powering on the system controller and disk enclosure, correct order,
 - disk arrangement in,
 - disk setup, [78](#)
 - disk status, [78](#)
 - RAID levels and default configuration, [55](#)
 - RAID10 initial setup, [61](#)
 - RAID50 capacity, [59](#)
 - RAID50 initial setup, [63](#)
- NSM 4150:RAID rebuild rate, [69](#)
- NTP
 - selecting, [120](#)
 - server, [120](#)
 - server, deleting, [121](#)
 - servers, changing list order, [121](#)



- off RAID
 - status, [71](#)
- ordering NTP access list, [121](#)

- overview
 - add-on applications, [317](#)
 - Centralized Management Console, [29](#)
 - clusters, [209](#)
 - disk replacement in special cases, [327](#)
 - Failover Manager, [189](#)
 - management groups, [171](#)
 - managers, [171](#)
 - network, [89](#)
 - provisioning storage, [221](#)
 - replacing a disk, [81](#)
 - reporting, [144](#)
 - setting date and time, [119](#)
 - SmartClone volumes, [263](#)
 - snapshots, [226](#), [245](#)
 - SNMP, [129](#)
 - storage category, [55](#)
 - volumes, [237](#)

P

- parity in RAID5, [58](#)
- Passwords
 - changing in Configuration Interface, [343](#)
- Pausing
 - monitoring, [313](#)
- pausing
 - scheduled snapshots, [256](#)
- performance
 - See I/O performance
- performance and iSCSI, [336](#)
- Performance Monitor
 - current SAN activity example, [298](#)
 - exporting data from, [315](#)
 - fault isolation example, [299](#)
 - learning about applications on the SAN, [299](#)
 - learning about SAN performance, [298](#)
 - load comparison of two clusters example, [301](#)
 - load comparison of two volumes example, [302](#)
 - NIC bonding example, [301](#)
 - overview, [297](#)
 - pausing, [313](#)
 - planning for SAN improvements, [300](#)
 - prerequisites, [297](#)
 - restarting, [313](#)
 - statistics, defined, [306](#)
 - understanding and using, [297](#)
 - workload characterization example, [298](#)

- Performance Monitor graph
 - changing, [313](#)
 - changing line color, [314](#)
 - changing line style, [314](#)
 - changing the scaling factor for, [315](#)
 - displaying a line, [314](#)
 - hiding, [314](#)
 - hiding a line, [314](#)
 - showing, [314](#)
- Performance Monitor window
 - accessing, [303](#)
 - graph, [305](#)
 - parts defined, [303](#)
 - saving to an image file, [316](#)
 - table, [306](#)
 - toolbar, [304](#)
- permanent variables, [138](#)
- Permissions
 - effect of levels, [293](#)
- permissions
 - administrative group, [127](#)
 - full, [126](#)
 - read modify, [126](#)
 - read only, [126](#)
- planning
 - data replication, [223](#)
 - RAID configuration, [67](#)
 - SmartClone volumes, [265](#)
 - snapshots, [226](#), [246](#)
 - volumes, [222](#), [237](#)
 - size, [221](#)
- planning capacity
 - full provisioning method, [222](#)
 - thin provisioning method, [222](#)
- point-in-time consistent snapshots
 - defined, [245](#)
- pool of storage, [171](#)
- powering off
 - disk, using CMC, [83](#), [84](#)
 - NSM 4150 system controller and disk enclosure, correct order, [47](#)
 - storage nodes, [48](#)
- powering on
 - disk, using CMC, [83](#), [85](#)
 - NSM 4150 system controller and disk enclosure, correct order, [47](#)
- preferred interface
 - in active-passive bond, [95](#)
 - in adaptive load balancing, [99](#)
 - in link aggregation dynamic mode bond, [97](#)
- preferred NTP server, [120](#)
- prerequisites, [263](#)
- Prerequisites for
 - assigning servers to volumes, [293](#)
 - servers, [290](#)

- prerequisites for
 - clusters, [209](#)
 - Performance Monitor, [297](#)
 - removing storage nodes from management group, [187](#)
 - snapshots, [245](#)
 - volumes
 - adding, [237](#)
 - deleting, [243](#), [244](#), [258](#), [261](#)
- primary interface, NICs, [116](#)
- primary volumes, [238](#)
- protocols, DHCP, [91](#)
- provisioning storage, [221](#)
 - and space allocation, [222](#)
 - best practices, [222](#)

Q

- quorum
 - and managers,
 - starting virtual manager to recover, [205](#)
 - stopping managers, [182](#)

R

RAID

- and data replication, [67](#)
- as data replication, [67](#)
- benefits, [57](#)
- changing RAID erases data, [70](#)
- configurations, [57](#)
- configurations defined, [57](#)
- configuring, [55](#), [56](#)
- default configuration on storage nodes, [55](#)
- definitions, [57](#)
- device, [59](#)
- device status, [56](#)
- levels and default configuration for:Dell 2950, [55](#)
- levels and default configuration for:NSM 160, [55](#)
- levels and default configuration for:NSM 2060, [55](#)
- levels and default configuration for:NSM 260, [55](#)
- levels and default configuration for:NSM 4150, [55](#)
- levels and default configuration:for DL380, [55](#)
- levels and default configuration:for HP LeftHand P4300, [55](#)
- levels and default configuration:for HP LeftHand P4500, [55](#)
- levels and default configuration:for IBM x3650, [55](#)
- managing, [56](#)
- Parity in RAID5, [58](#)
- planning configuration, [67](#)
- procedure for reconfiguring, [71](#)
- rebuild rate, [69](#)
- rebuilding, [86](#)
- reconfiguring, [70](#)
- reconfiguring requirements, [70](#)
- replacing a disk, [83](#), [85](#), [86](#)
- replication in a cluster, [68](#)
- requirements for configuring, [70](#)
- resyncing, [214](#)
- setting rebuild rate, [70](#)
- status, [71](#)
- status and data redundancy, [71](#)
- status and data transfer, [71](#)
- VSA, levels and default configuration for:VSA, [55](#)

- RAID (virtual), devices, [60](#)
- RAID levels defined, [327](#)
- RAID status
 - status, [71](#)

- RAID0
 - capacity, [57](#)
 - definition, [57](#)
 - devices, [60](#)
 - single disk replacement, [83](#)
- RAID1/10
 - single disk replacement, [83](#)
- RAID10
 - capacity, [57](#)
 - defined, [57](#)
- RAID5
 - capacity, [58](#)
 - configurations, [58](#)
 - defined, [58](#)
 - disk arrays, [63](#)
 - hot spares, [58](#)
- RAID5/50
 - :single disk replacement, [83](#)
- RAID50
 - capacity, [59](#)
 - defined, [59](#)
- RAID6
 - single disk replacement,
 - capacity, [59](#)
 - definition, [59](#)
 - devices, [66](#)
 - hot swapping, [59](#)
- raw storage, [233](#)
- re-create the RAID array ,
- read only permissions, [126](#)
- read only volumes, [250](#)
- read-modify permissions, [126](#)
- rebooting
 - storage nodes, [48](#)
- rebuild data
 - when not running manager, [328](#)
- rebuild volume data, [333](#)
- rebuilding
 - RAID, [86](#)
 - rate for RAID, [69](#)
- reconfiguring RAID, [70](#)
- recurring snapshots, [254](#)
- redundancy, data, [223](#)
- redundant array of independent disks
 - See RAID
- registering add-on applications, [320](#)
- registering features, [180](#)
 - Feature Registration tab, [320](#), [321](#)
 - for a management group, [180](#)
 - for a storage node, [50](#)
- registration information, [324](#)
- Remote Copy
 - backing out of evaluation, [318](#)
 - evaluating, [318](#)
 - registering, [320](#)
 - remote copies defined, [32](#)
 - remote volumes, [238](#)
- remote log files, [168](#)
 - adding, [168](#)
 - changing remote log file target computer, [169](#)
 - configuring target computer, [169](#)
 - removing old logs, [169](#)
- remote volumes, [238](#)
 - See Remote Copy User Manual
- Removing
 - statistics, [312](#)
- removing
 - administrative users from a Group, [127](#)
 - dedicated boot devices, [53](#)
 - DNS server, [114](#)
 - domain name from DNS suffixes list, [114](#)
 - ghost storage node after the data is rebuilt, [333](#)
 - monitored variables, [138](#)
 - old log files, [169](#)
 - SNMP trap recipient, [133](#)
 - storage nodes from cluster, [212](#)
 - storage nodes from management groups, [187](#)
 - prerequisites for, [187](#)
 - users from administrative groups, [127](#)
 - virtual manager, [207](#)
- Repair Storage Node
 - replacing a disk, [80](#)
- repair storage node, [216](#)
 - prerequisites, [216](#)
- repairing volumes
 - making unavailable priority volume available, [242](#)
- replacing
 - dedicated boot devices, [53](#)
 - disks, [80](#)
- replication
 - changing levels for volumes, [242](#)
 - data, [223](#)
 - level for volumes, [223](#)
 - levels allowed in clusters, [223](#)
 - RAID vs. volume replication, [67](#)
 - repairing storage node in cluster with replication, [216](#)
 - requirements for setting levels, [238](#)
- replication priority
 - and fault tolerance for volumes,
 - changing, [242](#)
 - changing priority, [242](#)
 - requirements for setting, [239](#)

- reporting overview, 144
- reports, 135
 - active, 135
 - details of Hardware Information report, 152
 - diagnostic, 145
 - disk, 72
 - disk setup for RAID, 73
 - generating, 150
 - hardware, 144
 - Hardware Information, 150
 - RAID setup, 59
 - saving Hardware Report to a file, 151
 - storage node statistics, 150
- requirements
 - system for Failover Manager on ESX Server, 190
- Requirements for
 - configuring CHAP for iSCSI, 291, 338
- requirements for
 - adding management group, 178
 - adding snapshots, 245
 - adding volumes, 238
 - application-managed snapshots, 248
 - changing SmartClone volumes, 283
 - changing volumes, 240
 - editing snapshots, 255
 - Failover Manager, 189
 - network interface bonding, 93
 - rolling back volumes, 258
 - scheduling snapshots, 254
 - snapshot schedules, 254
 - system for Failover Manager on VMware Server or Player, 189
 - using more than one Failover Manager, 190
 - virtual manager, 201
- resetting
 - DSM in Configuration Interface, 345
 - NSM to factory defaults, 345
- resolving host names, 44
- Restarting monitoring, 313
- restoring
 - management group, 184
 - storage node configuration files, 46
 - volumes, 257
- Restoring defaults for the Performance Monitor, 313
- restripping, volume, 84
- resuming scheduled snapshots, 256
- resyncing
 - RAID, 214
- return the repaired storage node to the cluster, 331
- Rolling Back a Volume
 - from application-managed snapshots, 259
- rolling back a volume, 257

- routing
 - adding network, 114
 - deleting, 115
 - editing network, 115
- routing tables
 - managing, 114

S

- safe to remove status, 81
- Sample interval, changing for the Performance Monitor, 309
- SAN
 - capacity of, 221
 - comparing the load of two clusters, 301, 308
 - comparing the load of two volumes, 302
 - current activity performance example, 298
 - determining if NIC bonding would improve performance, 301
 - fault isolation example, 299
 - learning about applications on the SAN, 299
 - learning about SAN performance, 298
 - monitoring performance, 297
 - planning for SAN improvements, 300
 - using Performance Monitor, 297
 - workload characterization example, 298
- SAN/iQ DSM for MPIO, 335
- saving
 - diagnostic reports, 145
 - history of one variable, 144
 - log files
 - of management group configurations, 184
 - log files for technical support, 167
 - log files of storage node configurations, 46
 - monitored variable log file, 143, 144
- Scaling factor
 - changing, 315
- scheduled snapshots, 254
 - pausing or resuming, 256
 - requirements for, 254
- scripting evaluation, 319
 - backing out of, 320
 - turning off, 319
 - turning on, 319
- searching for storage nodes, 39
- searching for storage nodes;, 29
- security
 - administrative, 171
 - of storage resources, 171
- selecting alerts to monitor, 136
- server access
 - SmartClone volumes, 267

- Servers
 - adding to management groups, [290](#)
 - assigning to volumes and snapshots, [292](#), [293](#)
 - deleting, [292](#)
 - editing, [291](#)
 - editing assignments to volumes and snapshots, [294](#)
 - prerequisites for, [290](#)
 - prerequisites for assigning to volumes, [293](#)
- servers
 - access to volumes and snapshots, [289](#)
 - defined, [32](#)
 - DNS
 - adding, [112](#)
 - editing IP or domain name, [113](#)
 - removing, [114](#)
 - iSNS
 - adding, [210](#)
 - NTP, [120](#)
 - editing, [120](#)
 - preferred, not preferred, [120](#)
- set ID LED, [45](#)
- setting
 - IP address, [91](#)
 - local bandwidth, [183](#)
 - RAID rebuild rate, [69](#)
- setting date and time, [119](#)
 - for management group, [119](#)
 - overview, [119](#)
 - procedure, [119](#), [121](#)
 - refreshing for management group, [119](#)
 - setting time zone, [122](#)
 - time zone on storage node, [119](#), [121](#), [122](#)
 - with NTP, [120](#)
 - without NTP, [121](#)
- setting up a RAID disk, [73](#)
- shared snapshots, [274](#)
- shutting down a management group, [185](#)
 - Management Group
 - shutting down, [184](#)
- single disk replacement in RAID0, [83](#)
- Single Host Configuration in iSCSI, [339](#)
- sites
 - defined, [32](#)
- size
 - changing for volumes, [242](#)
 - for snapshots, [227](#)
 - planning for
 - snapshots, [246](#)
 - volumes, [221](#)
 - requirements for volumes, [238](#)
- slow I/O, [214](#)
- SmartClone Volumes
 - making application-managed snapshot available after creating, [250](#), [251](#), [252](#)
- SmartClone volumes
 - assigning server access, characteristics of, shared versus individual, characteristics of, [267](#)
 - clone point, [272](#)
 - creating from application-managed snapshots, [260](#)
 - definition of, [263](#)
 - deleting, [284](#)
 - deleting multiple, [285](#)
 - editing, [284](#)
 - examples for using, [264](#)
 - glossary for, [264](#)
 - overview, [263](#)
 - planning, [265](#)
 - planning naming convention, [266](#)
 - planning space requirements, [265](#)
 - requirements for changing, [283](#)
 - uses for, [265](#)
 - utilization of, [282](#)
 - viewing with Map View, [280](#)
- SMTP
 - setting SMTP for alert notification, [142](#)
 - settings for alert notification, [142](#)
 - settings for alert notification, one variable, [142](#)
 - settings for alert notification, several variables, [143](#)
- Snapshots
 - assigning to servers, [292](#), [293](#)
 - editing server assignments, [294](#)

- snapshots
 - adding, [248](#)
 - adding schedules for, [255](#)
 - and upgrading software, [247](#)
 - application-managed, [245](#)
 - as opposed to Backups, [226](#)
 - changing thresholds, [250](#)
 - controlling server access to, [289](#)
 - copying a volume from, [250](#)
 - creating application-managed, [248](#)
 - creating application-managed for volume sets, [249](#)
 - defined, [32](#)
 - deleting, [261](#)
 - deleting schedules, [256](#)
 - editing, [250](#)
 - editing schedules, [255](#)
 - managing capacity
 - and scheduled snapshots, [254](#)
 - and thresholds, [226](#)
 - mounting, [250](#)
 - overview, [226](#), [245](#)
 - pausing or resuming, [256](#)
 - planning, [226](#), [246](#)
 - planning size, [246](#)
 - point-in-time consistent, [245](#)
 - prerequisites for, [245](#)
 - read/write
 - deleting temporary space, [253](#)
 - requirements for editing, [255](#)
 - rolling back a volume from, [257](#)
 - schedule requirements, [254](#)
 - scheduling, [254](#)
 - shared, [274](#)
 - size, [227](#)
 - temporary space for read/write snapshots, [235](#), [253](#)
 - using, [245](#)
 - versus backups, [245](#)
- SNMP
 - access control, [130](#)
 - Agents
 - community strings, [130](#)
 - disabling, [132](#)
 - agents
 - community strings, [130](#)
 - Clients, adding, [130](#)
 - MIB, [325](#)
 - overview, [129](#)
 - removing trap recipient, [133](#)
 - Traps
 - enabling, [133](#)
 - traps
 - disabling, [134](#)
 - editing recipient, [133](#)
 - using, [133](#)
 - using MIB, [132](#)
 - software
 - upgrading storage nodes, [49](#)
 - space allocation, [222](#)
 - space requirements
 - planning for SmartClone volumes, [265](#)
 - speed/duplex
 - configuring, [108](#)
 - VSA, [89](#)
 - Spoofing, [337](#)
 - starting
 - management group, [185](#)
 - managers on storage nodes, [181](#)
 - virtual manager to recover quorum, [205](#)
 - starting and stopping dedicated boot devices, [52](#)
 - startup and shutdown troubleshooting, [194](#)
 - static IP addresses and DNS, [112](#)
 - Statistics
 - adding, [310](#)
 - exporting performance to a CSV file, [315](#)
 - in the Performance Monitor defined, [306](#)
 - removing, [312](#)
 - viewing details of, [312](#)
 - statistics sample data
 - clearing, [312](#)
 - status
 - dedicated boot devices, [52](#)
 - NIC bond, [103](#)
 - RAID, [71](#)
 - safe to remove disk, [81](#)
 - storage node, [215](#)
 - storage server inoperable, [214](#)
 - storage server overloaded, [214](#)
 - stopping
 - managers, [181](#)
 - implications of, [182](#)
 - virtual manager, [207](#)

- storage
 - adding to a cluster, 212
 - configuration on storage nodes, 55
 - configuring, 55
 - overview, 55
- storage node
 - configuration categories, 43
 - expanded statistics detail, 150
 - logging in to an additional, 44
 - removing, 187
 - saving statistics to a file, 151
 - statistics, 150
- storage nodes
 - adding first one, 37, 178
 - adding to existing cluster, 211
 - adding to management group, 180
 - configuration file
 - backing up, 46
 - restoring, 46
 - configuration overview, 43
 - configuring, 37
 - configuring multiple, 40
 - default RAID configuration on, 55
 - details tab, 49
 - finding on network, 29, 37, 39
 - ghost storage node, 216
 - locating in a rack, 45
 - powering off, 48
 - rebooting, 48
 - registering, 50
 - removing from cluster, 212
 - removing from management group, 187
 - prerequisites for, 187
 - replacing disks, 80
 - repairing in clusters, 216
 - status of, 215
 - storage configuration of, 55
 - tasks, 43
 - upgrading software, 49
- storage pool, 171
- storage server inoperable, 214
- storage server overloaded, 214
- storage server status and VSA, 214
- storage space
 - raw space, 233
- storage, provisioning, 221
- Subscriber's Choice, HP, 27
- synchronizing time on storage nodes, 254
- System Management homepage, logging into, 130
- system requirements
 - for Failover Manager on ESX Server,
 - for Failover Manager on VMware Server or Player,

T

- Tab window, 33
- TCP
 - speed and duplex,
 - frame size, 110
 - status, 107
 - status tab, 107
- TCP/IP tab, 90
- technical support
 - HP, 27
 - saving log files
 - for, 167
 - service locator website, 27
- temporary space
 - deleting, 253
 - for read/write snapshots, 235, 253
 - making application-managed snapshot
 - available after converting, 250, 251, 252
- thresholds
 - capacity management and snapshot, 226
 - changing
 - for a snapshot, 250
 - requirements for
 - changing in snapshots, 255
- time
 - editing NTP server, 120
 - NTP servers, preferred, not preferred, 120
 - selecting NTP, 120
 - setting
 - with NTP, 120
 - without NTP, 121
 - synchronizing on storage nodes, 254
 - zone, setting on storage node, 119, 121, 122
- time remaining on evaluation period, 318
- Time zone
 - changing for the Performance Monitor, 309
- time zone
 - setting, 122
- Toolbar
 - Performance Monitor window, 304
- toolbar
 - SmartClone map view, 280
- trap recipient
 - removing, 133
- traps
 - disabling SNMP, 134
 - editing SNMP recipient, 133
 - enabling SNMP, 133
 - SNMP, 133
- troubleshooting
 - network settings to find Failover Manager, 195
 - startup and shutdown options, 194

- troubleshooting clusters
 - repair storage node, [214](#)
 - slow I/O, [214](#)
- type
 - See volumes
- U**
- updating
 - hardware information report, [150](#)
 - manger IP addresses, [117](#)
- upgrading
 - storage node software,
- upgrading software
 - copying to the storage node, [50](#)
 - copying upgrade files, [49](#)
- user
 - adding a group to a user, [124](#)
 - administrative, [123](#)
 - administrative default user, [123](#)
 - changing a user name, [124](#)
 - deleting administrative, [124](#)
 - editing, [124](#)
 - password, [124](#)
- utilization
 - of clone points and SmartClone volumes, [282](#)

- V**
- variables, monitored
 - adding, [136](#)
 - downloading log file for, [143](#), [144](#)
 - editing, [137](#)
 - list of ,
 - permanent, [138](#)
 - removing, [138](#)
 - viewing summary of, [141](#)
- verifying NIC bond, [102](#)
- VI Client
 - recreating disk for VSA, [77](#)
- Viewing
 - statistics details, [312](#)
- viewing
 - disk report, [73](#)
 - disk setup reportdisk, [72](#)
 - monitored variable summary, [141](#)
- viewing clone points, volumes and snapshots, [282](#)
- viewing SmartClone volumes, [279](#)
- viewing:RAID setup report, [59](#)
- Virtual IP Address
 - host storage node, [335](#)

- virtual IP address, [335](#)
 - and iSCSI, [335](#)
 - changing, iSCSI, [212](#)
 - configuring for iSCSI for, [210](#)
 - gateway session when using load balancing, [336](#)
 - removing, iSCSI volume, [212](#)
- virtual machine, [173](#)
- virtual manager
 - adding, [204](#)
 - benefits of, [201](#)
 - configurations for using, [200](#)
 - configuring, [204](#)
 - function of, [200](#)
 - overview, [200](#)
 - removing, [207](#)
 - starting to recover quorum, [205](#)
 - stopping, [207](#)
- virtual RAID
 - data safety and availability, [69](#)
- virtual storage node
 - data replication, [67](#)
 - RAID device, [60](#)
- VMware ESX Server, [60](#)
- VMware Server, [173](#)
- volume availability, [214](#)
- volume sets
 - creating application-managed snapshots for, [249](#)
 - deleting application-managed snapshots for, [261](#)
- volume size
 - best practice for setting, [222](#)
- Volumes
 - assigning to servers, [292](#), [293](#)
 - comparing the load of two, [302](#)
 - editing server assignments, [294](#)
 - formatting for use, [295](#)
 - iSCSI, [337](#)
 - and CHAP, [337](#)
 - logging on to, [295](#)
 - setting as persistent targets, [295](#)

volumes

- Access Volume wizard, [39](#)
- adding, [239](#)
- changing
 - clusters, [242](#)
 - descriptions, [241](#)
 - replication levels, [242](#)
- changing size, [242](#)
- controlling server access to, [289](#)
- creating SmartClone, [276](#)
- creating using the Wizard, [38](#)
- defined, [32](#)
- deleting, [243](#)
- editing, [240](#)
- making unavailable priority volume available, [242](#)
- mounting file systems on, [237](#)
- overview, [237](#)
- planning, [222](#), [237](#)
 - size, [221](#)
 - type, [238](#)
- prerequisites for
 - adding, [237](#)
 - deleting, [243](#), [244](#), [258](#), [261](#)
- replication level, [223](#)
- replication priority, [225](#)
- requirements for
 - adding, [238](#)
 - changing, [240](#)
- restripping, [84](#)
- rolling back, [257](#)
 - application server requirements, [258](#)
- type
 - primary, [238](#)
 - remote, [238](#)

volumes and snapshots

- availability, [51](#)

VSA

- cloning, [178](#)
- disk status, [77](#)
- frame size, [89](#)
- hardware diagnostics, [149](#)
- hardware report, [152](#)
- network interface, [89](#)
- NIC bonding, [89](#)
- NIC flow control, [89](#)
- RAID levels and default configuration, [55](#)
- RAID rebuild rate, [69](#)
- reconfiguring RAID, [70](#)
- recreate disk, [77](#)
- speed/duplex, [89](#)
- storage server overloaded, [214](#)
- virtual RAID and data safety and availability, [69](#)

W

warnings

- :check Safe to Remove status,
- all storage nodes in a cluster operate at a capacity equal to that of the smallest capacity, [209](#)
- changing RAID erases data, [70](#)
- cloning VSA, [178](#)
- deleting management group causes data loss, [187](#)
- DHCP
 - static IP addresses, [92](#)
 - unicast communication, [92](#)
- return repaired node to same place, [217](#)
- stopping manager can cause data loss, [182](#)

websites

- HP ,
- HP Subscriber's Choice for Business, [27](#)
- product manuals, [27](#)

windows

- Alert window, [30](#)
- navigation window, [30](#)
- Tabs window, [30](#)

wizards

- Getting Started Launch Pad, [32](#)

write failure, warnings

- write failure, [222](#)

