

# HP SAN Design Reference Guide

## **Abstract**

This reference document provides information about HP SAN architecture, including Fibre Channel, iSCSI, FCoE, SAN extension, and hardware interoperability. Storage architects and system administrators can use this document to plan, design, and maintain an HP SAN.



© Copyright 2001-2011 Hewlett-Packard Development Company, L.P.

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

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

### **Acknowledgments**

Itanium® is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle® and Java® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

### **Warranty**

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

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

---

# Contents

I Architecture.....	15
1 SAN design overview.....	16
SAN solutions.....	16
HP SAN implementations.....	17
SAN components.....	18
Fibre Channel technology.....	18
Storage area networks.....	18
SAN infrastructure.....	19
Fabrics.....	19
SAN scaling.....	19
Fibre Channel switches.....	19
SAN design approaches.....	20
SAN design considerations.....	20
2 SAN fabric topologies.....	23
Fabric topologies.....	23
Routed SAN fabrics.....	24
FCoE SAN fabrics.....	24
Benefits.....	24
Single-switch fabric.....	24
Switch models.....	24
Benefits.....	25
Cascaded fabric.....	25
Switch models.....	26
Benefits.....	26
Meshed fabric.....	26
Switch models.....	27
Benefits.....	27
Ring fabric.....	27
Switch models.....	29
Benefits.....	29
Core-edge fabric.....	29
Core-edge fabric types.....	30
Switch models.....	32
Benefits.....	32
Topology data access.....	32
Topology maximums.....	33
B-series switches.....	33
C-series switches.....	34
H-series switches.....	34
M-series switches.....	34
Routed fabric topologies.....	35
B-series Meta SAN.....	35
B-series Virtual Fabrics with IFR.....	36
C-series VSANs with IVR.....	37
H-series switches with TR.....	37
FCoE fabric topologies.....	38
Data availability.....	38
Levels.....	38
Considerations.....	41
Topology migration.....	41

Nondisruptive migration.....	41
Migrating a cascaded fabric SAN.....	42
Migrating a meshed fabric SAN.....	42
Migrating a ring fabric SAN.....	42
<b>3 Fibre Channel routing.....</b>	<b>43</b>
Fibre Channel routing overview.....	43
Fabric, Virtual Fabric, and VSAN independence.....	44
Fabric services.....	44
Worldwide Name.....	44
Import and export.....	44
Routing table.....	44
SAN scaling and routing.....	44
Switch scaling.....	44
Scaling by routing.....	46
Fibre Channel routing implementations.....	46
Fibre Channel routing techniques.....	46
B-series fabric groups.....	47
B-series fabric partitioning using Virtual Fabrics.....	48
B-series Virtual Fabrics architecture.....	48
C-series fabric partitioning.....	48
H-series switch fabric routing.....	48
B-series, C-series, and H-series routing differences.....	48
Fabric redundancy and routing.....	51
Supported routing configurations.....	52
Routing and core-edge fabrics.....	52
Routing through an IP network.....	53
High-availability router configurations.....	55
400 MP Router and MP Router Blade use cases.....	56
H-series switches with TR configurations.....	56
Routing use cases.....	57
<b>4 Fibre Channel over Ethernet.....</b>	<b>61</b>
HP FCoE solutions overview.....	61
FCoE technology.....	62
Converged network switch ports.....	62
Interoperability.....	62
HP FCoE products.....	63
Converged network switches and blades.....	63
HP A5820 Ethernet/FCoE Switch.....	63
HP 2408 FCoE Converged Network Switch .....	64
HP StorageWorks DC SAN Director Switch 10/24 FCoE Blade.....	64
HP C-series Nexus 5010/5020 Converged Network Switches.....	65
Converged network adapters .....	65
HP CN1000E and CN1100E.....	66
CN1000Q.....	66
FCoE storage systems.....	67
FCoE configuration rules.....	69
Server support.....	74
HP A5820 Converged Network Switch.....	74
Usage.....	74
FCoE Converged Network Switch fabric rules.....	74
Firmware, CNA, operating systems, and storage products support.....	75
B-series FCoE Converged Network switch.....	75
Features.....	76
Operating systems and storage products.....	76

C-series FCoE Converged Network switches.....	76
C-series Nexus 5020 FCoE Converged Network Switch.....	77
C-series Nexus 5010 FCoE Converged Network Switch.....	77
Expansion modules for the C-series Nexus 5000 Series switches.....	77
Features.....	78
Usage.....	78
FCoE Converged Network switch fabric rules.....	78
Operating systems and storage products.....	78
Storage system support.....	78
Operating system support.....	79
Boot from SAN support.....	79
Multipathing software support.....	81
<b>II Fabric infrastructure rules.....</b>	<b>82</b>
<b>5 B-series switches and fabric rules.....</b>	<b>83</b>
B-series Fibre Channel switches, 1606 Extension SAN Switches, and MP Routers.....	83
Model numbering.....	84
Model naming.....	84
Switch models.....	84
Features.....	88
Usage.....	93
Fibre Channel switch fabric rules.....	94
Operating systems and storage products.....	95
B-series Fibre Channel switch and fabric rules.....	95
B-series Encryption Switch fabric rules.....	99
1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules.....	100
Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules.....	103
MP Router fabric rules.....	107
Core switch addressing mode.....	112
Zoning limits and enforcement.....	112
Zoning guidelines for B-series Fibre Channel switches.....	113
<b>6 C-series switches and fabric rules.....</b>	<b>115</b>
C-series Fibre Channel switches.....	115
Model naming.....	117
Switch models.....	117
Features.....	123
Usage.....	125
Fibre Channel switch fabric rules.....	126
Operating systems and storage products.....	126
Fabric rules for C-series Fibre Channel switches.....	127
ISL maximums.....	128
Zoning limits and enforcement.....	129
C-series VSAN high availability.....	129
<b>7 H-series switches and fabric rules.....</b>	<b>131</b>
H-series switches.....	131
Model numbering.....	132
Model naming.....	132
Switch models.....	132
Features.....	132
Usage.....	134
Fabric rules.....	134
Servers, operating systems, and storage products.....	135
Fabric rules for H-series switches.....	135
ISL maximums.....	136

Fabric rules for H-series switches with TR.....	136
Zoning limits and enforcement.....	138
<b>8 M-series switches and fabric rules.....</b>	<b>140</b>
M-series switches.....	140
Model naming.....	140
Switch models.....	140
Features.....	142
Usage.....	142
Fibre Channel fabric rules.....	142
Operating systems and storage products.....	142
Fabric rules for M-series switches.....	143
ISL maximums.....	144
Zoning limits and enforcement.....	144
<b>9 SAN fabric connectivity and switch interoperability rules.....</b>	<b>146</b>
SAN fabric connectivity rules.....	146
Switch port interfaces.....	146
Device port interfaces.....	146
Fiber optic cables.....	146
Fiber optic cable loss budgets.....	148
Storage product interface, switches, and transport distance rules.....	150
SAN fabric switch interoperability rules.....	154
Interoperable routed fabric.....	155
Dual interoperable, heterogeneous SAN fabrics.....	155
Interoperable, heterogeneous switch fabrics.....	155
Third-party switch support.....	156
SAN performance considerations.....	156
Infrastructure factors.....	156
Performance guidelines.....	157
<b>III Host and storage system rules.....</b>	<b>158</b>
<b>10 Heterogeneous server rules.....</b>	<b>159</b>
SAN platform rules.....	160
Heterogeneous storage system support.....	160
Blade servers.....	161
HP c-Class BladeSystem Virtual Connect.....	161
HP Virtual Connect advantages.....	162
HP VC-FC connectivity guidelines.....	162
BladeSystem with Brocade Access Gateway mode.....	163
Failover policy and failback policy.....	164
AG mode considerations.....	164
AG mode connectivity guidelines.....	165
BladeSystem with Cisco N_Port Virtualization mode.....	165
Failover policy.....	166
NPV mode considerations.....	167
NPV mode connectivity guidelines.....	167
NPV with FlexAttach.....	168
HP BladeSystem c3000 enclosure considerations.....	169
HBA N_PortID Virtualization.....	170
HBA NPIV considerations.....	170
HBA NPIV connectivity guidelines.....	170
NonStop servers (XP only).....	171
HP-UX SAN rules.....	183
HP OpenVMS SAN rules.....	185
Host-based volume shadowing.....	186

HP Tru64 UNIX SAN rules.....	187
Apple Mac OS X SAN rules.....	188
IBM AIX SAN rules.....	190
Linux SAN rules.....	191
Linux multipathing coexistence support.....	193
Microsoft Windows SAN rules.....	193
Sun Solaris SAN rules.....	196
VMware ESX SAN rules.....	198
VMware ESX multipathing coexistence support.....	199
Citrix Xen SAN rules.....	199
Heterogeneous SAN storage system coexistence.....	200
Common SAN coexistence.....	200
Common server access.....	201
Common server, different HBAs.....	201
Common server, common HBA.....	201
Server zoning rules.....	202
<b>11 P2000 and MSA storage system rules.....</b>	<b>203</b>
HP StorageWorks P2000 G3 FC/iSCSI, Modular Smart Array 2000fc G2, and Modular Smart Array 2000fc.....	203
Heterogeneous SAN support.....	203
Configuration rules.....	204
Zoning.....	205
Maximums.....	205
Management software support.....	206
Tape storage.....	206
HP StorageWorks Modular Smart Array 1500.....	207
Homogeneous SAN support.....	207
Heterogeneous SAN support.....	207
Configuration rules.....	208
Zoning.....	208
Maximums.....	208
Tape storage.....	209
HP StorageWorks Modular Smart Array 1000.....	209
Homogeneous SAN support.....	209
Heterogeneous SAN support.....	210
Configuration rules.....	210
Zoning.....	211
Maximums.....	211
Tape storage.....	212
<b>12 P6000/EVA storage system rules.....</b>	<b>213</b>
P6000/EVA storage.....	213
Heterogeneous SAN support.....	214
Configuration rules.....	214
Configuration parameters.....	216
P6000/EVA data migration.....	217
MPX200 Multifunction Router with data migration.....	218
Data migration considerations.....	219
HP P6000 Continuous Access SAN integration.....	220
Zoning.....	221
Tape storage.....	221
P6000/EVA SAN boot support.....	221
Storage management server integration.....	221
Cabling.....	222
Level 4 NSPOF configuration.....	222

Dual-channel HBA configurations.....	225
<b>13 P9000/XP and VA storage system rules.....</b>	<b>227</b>
P9000/XP and VA storage systems.....	227
Heterogeneous SAN support.....	227
Configuration rules.....	228
Zoning.....	229
Tape storage.....	229
P9000/XP and VA SAN boot support.....	230
LUN Configuration and Security Manager XP support.....	231
P9000/XP data migration.....	231
MPX200 Multifunction Router with data migration.....	231
Data migration considerations.....	233
<b>14 SVSP storage system rules.....</b>	<b>234</b>
SVSP storage systems.....	234
Features.....	234
Software.....	235
Hardware for a single SVSP domain.....	235
Heterogeneous SAN support.....	236
Configuration rules.....	236
Storage rules.....	237
SVSP supported third-party arrays.....	238
SVSP data migration.....	238
MPX200 Multifunction Router with data migration.....	238
SVSP Continuous Access.....	238
Configuration parameters.....	239
SVSP best practices.....	240
SVSP SAN boot support.....	240
Storage management server integration.....	240
<b>15 HP 3PAR storage system rules.....</b>	<b>241</b>
HP 3PAR storage systems.....	241
Heterogeneous SAN support.....	241
Configuration rules.....	243
Configuration parameters.....	243
HP 3PAR data migration.....	243
MPX200 Multifunction Router with data migration.....	243
Zoning.....	246
HP 3PAR SAN boot support.....	246
HP 3PAR storage management.....	246
<b>16 Enterprise Backup Solution.....</b>	<b>247</b>
<b>IV SAN extension and bridging.....</b>	<b>248</b>
<b>17 SAN extension.....</b>	<b>249</b>
SAN extension overview.....	249
SAN extension technology.....	250
SAN-iSCSI bridging technology.....	250
Fibre Channel long-distance technology.....	251
Fiber optic transceivers.....	251
Wavelength division multiplexing.....	253
WDM overview.....	253
WDM network implementation.....	253
WDM system architectures.....	253
WDM system characteristics.....	254
HP coarse wave division multiplexing.....	254



Third-party WDM products.....	255
Extended fabric settings for Fibre Channel switches.....	256
B-series switch settings.....	256
B-series trunking and WDM.....	257
C-series switch settings.....	257
H-series switch settings.....	257
M-series switch settings.....	258
Multi-protocol long-distance technology.....	259
Fibre Channel over Internet Protocol.....	259
FCIP mechanisms.....	260
FCIP link configurations.....	260
FCIP single-link configuration.....	260
FCIP dual-link configuration.....	260
FCIP shared-link configuration.....	260
FCIP network considerations.....	261
FCIP bandwidth considerations.....	261
Recommendations for managing bandwidth with FCIP.....	261
FCIP gateways.....	263
Third-party QoS and data encryption FCIP products.....	263
FCIP QoS products.....	263
WAN accelerator products.....	264
Fibre Channel over SONET.....	264
FC-SONET IP link configurations.....	264
FC-SONET dual-link configuration.....	264
FC-SONET shared-link configuration.....	264
FC-SONET network considerations.....	265
Third-party SONET gateways.....	265
Fibre Channel over ATM.....	265
HP multi-protocol long-distance products.....	265
HP SAN extension products summary and usage.....	265
HP IP Distance Gateway (mpx110).....	266
IP Distance Gateway configuration examples.....	267
Configuration rules.....	274
General configuration rules.....	274
Operating system and multipath support.....	275
EVA storage system rules.....	275
XP storage system rules.....	275
XP storage system software.....	275
Fibre Channel switch and firmware support.....	275
MPX200 Multifunction Router with FCIP.....	276
MPX200 Multifunction Router FCIP configuration examples.....	276
FCIP Configuration rules.....	279
General FCIP configuration rules.....	279
Operating system and multipath support.....	279
Storage system rules.....	279
B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade.....	280
1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements.....	281
1606 Extension SAN Switch configuration examples.....	281
B-series 400 MP Router and MP Router Blade.....	283
400 MP Router and MP Router Blade features and requirements.....	283
400 MP Router and MP Router Blade configuration examples.....	284
B-series MP Router.....	285
B-series MP Router features and requirements.....	285
B-series MP Router configuration examples.....	286

C-series MDS 9216i, 9222i, IPS-4, IPS-8, 14/2 Multiprotocol Services Modules, 18/4 Multiservice Modules.....	287
HP storage replication products.....	288
SAN extension best practices for HP P6000 Continuous Access.....	289
HP P6000 Continuous Access with XCS 10x or XCS 09x.....	289
HP P6000 Continuous Access with XCS 6.x.....	291
HP P6000 Continuous Access with VCS 4.x.....	293
HP P9000 (XP) Continuous Access.....	296
OpenVMS host-based volume shadowing.....	310
Certified third-party WDM, iFCP, and SONET products.....	311
Certified third-party WDM products.....	311
Certified third-party iFCP and SONET products.....	311
<b>18 iSCSI storage.....</b>	<b>312</b>
iSCSI overview.....	312
iSCSI and Fibre Channel.....	312
iSCSI concepts.....	312
Initiator and target devices.....	313
iSCSI naming.....	313
Discovery mechanisms.....	313
Service Location Protocol .....	314
Static configuration.....	314
SendTargets command.....	314
Internet Storage Name Service.....	314
Sessions and logins.....	314
Security.....	315
Software and hardware iSCSI initiators.....	315
iSCSI boot.....	315
iSCSI storage network requirements.....	316
HP Native iSCSI products.....	316
P6300/P6500 EVA.....	316
P6300/P6500 EVA overview.....	316
P6300/6500 EVA iSCSI support.....	316
HP StorageWorks Modular Smart Array 1510i.....	317
MSA1510i overview.....	317
MSA1510i support.....	317
Network device support.....	317
Server support.....	317
Operating system support.....	317
Path failover software.....	318
Management software support.....	318
MSA1510i rules.....	318
HP StorageWorks MSA/2000 family of iSCSI SAN arrays.....	318
MSA2000i G2 and MSA2000i overview.....	318
P2000 G3 FC/iSCSI, P2000 G3 10Gb iSCSI, P2000 G3 iSCSI overview.....	318
P2000 G3 FC/iSCSI, MSA2000i G2 and MSA2000i support.....	319
Server support.....	319
Operating system support.....	320
Path failover software.....	320
Management software support.....	320
Maximum configurations.....	321
HP StorageWorks P4000 SAN.....	321
P4000 SAN Solutions overview.....	321
P4000 SAN Solutions support.....	322
Path failover software.....	322

Management software support.....	322
Maximum configurations.....	322
HP iSCSI bridge products.....	323
Bridging and routing.....	323
iSCSI bridge to Fibre Channel.....	323
MPX200 Multifunction Router with iSCSI for EVA storage.....	324
MPX200 simultaneous operation.....	324
MPX200 configuration options.....	324
MPX200 iSCSI rules and supported maximums .....	328
P6000/EVA storage system rules and guidelines .....	329
HP P6000 Command View and MPX200 management rules and guidelines.....	329
EVA storage system software.....	330
Fibre Channel switch and fabric support.....	330
Operating system and multipath software support.....	330
MPX200 Multifunction Router with iSCSI for XP storage.....	330
MPX200 configuration options.....	331
MPX200 iSCSI rules and supported maximums.....	332
XP storage system rules and guidelines.....	332
Operating system and multipath software support.....	332
EVA and EVA4400 iSCSI Connectivity Option .....	333
Hardware support.....	335
mpx100/100b data transport.....	335
Fibre Channel switches.....	336
Storage systems.....	336
Software support.....	336
Management software.....	336
Multipath software.....	337
HP P6000 Continuous Access.....	338
Operating systems and network interface cards.....	339
NIC Teaming.....	339
iSCSI initiators.....	339
iSCSI boot.....	339
EVA and EVA4400 iSCSI Connectivity Options supported maximums.....	340
General rules for the EVA and EVA4400 iSCSI Connectivity Options.....	340
B-series iSCSI Director Blade.....	341
Blade overview.....	341
Hardware support.....	342
Storage systems.....	342
Fibre Channel switches.....	342
Software support.....	343
Operating systems and network interface controllers.....	343
Network Teaming.....	343
B-series management applications.....	343
iSCSI initiators.....	343
Scalability rules.....	343
B-series MP Router and iSCSI.....	344
Hardware support.....	344
Fibre Channel switches.....	344
Storage systems.....	344
Software support.....	344
Operating systems and network interface controllers.....	344
Network Teaming.....	344
iSCSI initiators.....	345
Configuration rules.....	345
C-series iSCSI.....	345

Modules overview.....	345
Hardware support.....	346
Storage systems.....	346
Fibre Channel switches.....	347
Software support.....	347
Operating systems and network interface controllers.....	347
HP Network Teaming.....	347
C-series management applications.....	348
iSCSI initiators.....	348
Configuration rules.....	348
HP ProLiant Storage Server iSCSI Feature Pack.....	349
Overview.....	349
HP ProLiant Storage Server iSCSI Feature Pack support.....	349
Hardware support.....	349
Application support.....	350
Management software support.....	350
iSCSI Initiator support rules.....	350
HP ProLiant Storage Server iSCSI license upgrade options.....	350
Snapshot.....	350
Clustering.....	351
Direct Backup.....	351
Designing a Microsoft Exchange solution with iSCSI Feature Pack.....	351
Network design.....	351
Hardware selection.....	351
Exchange storage design.....	352
Supported load with Exchange.....	352
Sample iSCSI NAS Microsoft Exchange Server 2003 configuration.....	353
HP ProLiant DL380 G4 Storage Server configuration.....	353
<b>19 Network Attached Storage.....</b>	<b>354</b>
Storage server and SAN integration.....	354
Configuration and zoning rules.....	354
HP ProLiant Storage Servers.....	354
Features.....	355
Software.....	355
Hardware.....	355
Fabric rules.....	356
Storage rules.....	356
<b>V Storage security, best practices, and support information.....</b>	<b>357</b>
<b>20 Storage security.....</b>	<b>358</b>
Storage security threats.....	358
Storage security compliance.....	359
Security technologies.....	359
IP SAN security technologies.....	359
Fibre Channel SAN security technologies.....	360
Encryption security technologies.....	360
Key management.....	361
Organizational security policies.....	361
HP security strategy.....	361
HP Secure Advantage.....	361
Resource protection.....	361
Data protection.....	362
Security validation.....	363
Storage security best practices.....	363

Assessing security risks.....	363
Managing organizational risks.....	363
Data security implementations.....	363
HP storage security solutions.....	364
C-series Storage Media Encryption.....	364
C-series SAN-OS security.....	365
C-series IP SAN security.....	366
B-series Encryption Switch and Encryption FC Blade security.....	366
B-series Fabric OS security.....	367
Resource protection.....	367
Data protection.....	369
Security validation.....	370
Key management.....	370
<b>21 Best practices.....</b>	<b>372</b>
SAN planning.....	372
Design specification.....	372
SAN topology.....	373
Multi-fabric SANs.....	373
SAN and fabric monitoring.....	373
Failover protection.....	373
Data access patterns.....	374
ISL ratio.....	374
Incremental SAN expansion.....	374
SAN configuration .....	374
Fibre Channel switch configuration.....	375
Server setup.....	375
Storage system configuration.....	375
Storage-based LUN masking.....	375
Zoning.....	376
Zoning enforcement.....	376
Zoning guidelines.....	376
Zoning by operating system.....	377
Zoning by HBA.....	377
Zoning by application.....	377
Zoning by port allocation.....	377
EBS zoning.....	377
Zone naming.....	377
Naming by identifier type.....	378
Case sensitivity of fabric identifiers.....	378
Server naming.....	378
Storage system naming.....	378
FCoE switch configuration quick-setup instructions.....	379
HP A5820 FCoE Converged Network Switch quick setup.....	379
2408 FCoE Converged Network Switch and DC SAN Director Switch 10/24 FCoE Blade quick setup.....	381
HP C-series Nexus 5010/20 Converged Network Switch quick setup.....	383
SAN scaling.....	389
Cascaded fabric expansion.....	389
Meshed fabric expansion.....	390
Ring fabric expansion.....	390
Core-edge fabric expansion.....	390
SAN fabric merging.....	390
Fabric segmentation errors.....	391
Switch configuration parameters.....	391

Independent fabric merge.....	391
High-availability redundant fabric merge.....	391
<b>22 Support and other resources.....</b>	<b>393</b>
Contacting HP.....	393
HP technical support.....	393
Subscription service.....	393
Documentation feedback.....	393
New and changed information in this edition.....	393
Related information.....	393
Typographic conventions.....	394
Customer self repair.....	395
<b>Glossary.....</b>	<b>396</b>
<b>Index.....</b>	<b>404</b>

---

# Part I Architecture

SAN architecture is presented in these chapters:

- “SAN design overview” (page 16)
- “SAN fabric topologies” (page 23)
- “Fibre Channel routing” (page 43)
- “Fibre Channel over Ethernet” (page 61)

---

# 1 SAN design overview

SANs provide the data communication infrastructure for advanced, cost-efficient storage systems. SAN technology offers investment protection, management features, and I/O price performance to minimize capital expense. HP SAN architecture provides open network storage solutions for all sizes and types of businesses, including small-to-medium-sized IT departments and enterprise environments.

This chapter describes the following topics:

- “SAN solutions” (page 16)
- “HP SAN implementations” (page 17)
- “SAN components” (page 18)
- “Fibre Channel technology” (page 18)
- “Storage area networks” (page 18)
- “SAN infrastructure” (page 19)
- “Fibre Channel switches” (page 19)
- “SAN design approaches” (page 20)
- “SAN design considerations” (page 20)

## SAN solutions

SANs provide flexibility in system management, configuration, connectivity, and performance to meet the needs of the changing business environment. For the most challenging IT problems, SANs offer resilient solutions:

- **Open systems**  
SANs support various operating systems and servers to meet your operational requirements. A robust storage infrastructure accommodates new business models, unexpected growth, and corporate reorganizations.
- **Fast backup and restore**  
SANs remove backup and recovery traffic from the LAN, reducing congestion, improving backup windows, and efficiently utilizing storage resources. You can use centrally managed, high-performance tape libraries to reduce backup overhead.
- **Business continuance**  
SANs can eliminate single points of failure, incorporate failover software, and support mirroring at geographically dispersed data centers for disaster recovery. You can quickly restore productivity after a power failure or component downtime.
- **High availability**  
Redundant fabric designs, storage replication, dynamic failover protection, traffic rerouting, and server clustering enable SANs to provide enterprise-class availability to open systems servers.
- **Server and storage consolidation**  
Multiple servers and backup systems can share storage for efficient processing and increased availability.
- **Cost savings**  
SAN total cost of ownership is typically less than DAS. The business realizes a higher return on investment because sharing storage among servers utilizes capacity more efficiently, and



expenses for backup hardware are reduced. Increased system availability can help prevent costly downtime and lost data.

- **Centralized management**

You can manage consolidated storage by using web-based tools from any location, thus reducing labor costs.

- **Security**

SANs support network security measures, such as authentication, authorization, access control, and zoning.

- **Online scalability**

You can add storage capacity or expand the fabric as needs change. You can add and remove servers, and increase, change, or reassign storage while the SAN is online.

- **Modularity**

Modular design simplifies SAN scalability and increases ROI by consolidating and sharing systems.

Your SAN can incorporate all of these features, or you can start with a small SAN and add features as your business needs change.

## HP SAN implementations

You can configure a custom SAN by choosing components and following the HP design rules. HP SAN designs employ a configuration philosophy that supports comprehensive SAN implementations.

- **Flexible design and deployment**

HP provides standard topologies and design rules to meet the widest range of requirements for small office environments, mid-range business systems, and enterprise-class installations. The design rules and methods described in this guide enable change and expansion as needs arise.

- **Incremental scaling**

HP SANs maximize value by optimizing features and functionality of the SAN components. You can expand your SAN over time by adding capacity and features as required.

- **Interoperability**

HP SAN designs support multiple operating system, server, storage system, and SAN infrastructure component types.

- **Geographically dispersed installations**

HP provides components to meet local and long-distance connectivity requirements.

For information about:

- SAN infrastructure solutions, see the Storage Networking website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

- SAN deployment in small-to-medium-sized businesses, see the HP Servers & Storage expertise center website:

<http://www.hp.com/sbso/serverstorage>

- The Small Business SAN Kit, see the MSA1000 website:

<http://h18006.www1.hp.com/products/storageworks/msa1000smb/index.html>

## SAN components

A SAN consists of the following hardware and software components:

- **Switches**  
A Fibre Channel switch creates the fabric of the SAN. By interconnecting switches, you can create scalable SANs with thousands of port connections.
- **Routers, bridges, and gateways**  
Router functionality provides high levels of scalability, dynamic device sharing, and Fibre Channel network fault isolation. Routers, bridges, and gateways extend the SAN over long distances and enable integration of multi-protocol technologies.
- **Storage devices**  
A SAN can integrate multiple storage system types, such as disk arrays and tape libraries, to allocate storage efficiently.
- **Servers and HBAs**  
HBAs connect the server to the SAN. HBA drivers provide an intelligent interface to the switches and minimize CPU overhead.
- **Cabling and cable connectors**  
Fiber optic cables provide the physical connections between SAN components.
- **SAN management applications**  
HP applications manage and monitor components and ensure optimal SAN operation.

## Fibre Channel technology

Fibre Channel is a comprehensive set of standards for communication among servers, storage systems, and peripheral devices. A Fibre Channel network provides connectivity among heterogeneous devices and supports multiple interconnect topologies.

The network can be connected to a variety of storage systems:

- RAID arrays
- Tape devices and backup libraries

Fibre Channel technology supports simultaneous use of these transport protocols:

- IP
- SCSI
- iSCSI

For the latest information on Fibre Channel and related technologies, see the following website:

<http://www.incits.org>

## Storage area networks

General-purpose networks, such as LANs, enable communication between servers. A SAN uses multiple paths to connect servers and storage systems. To take full advantage of its capabilities, the SAN is maintained separately from parallel general-purpose networks.

The network topology is the physical arrangement of interconnected hardware components. In a basic topology, a Fibre Channel switch interconnects multiple servers and a storage system. To protect against hardware failure, high-availability topologies connect redundant systems. You can connect a complex and extensible network across long distances by choosing the required topology and appropriate components, and then connecting devices with fiber optic cable.

## SAN infrastructure

You use fabric switches to create the SAN communication paths. The number of storage systems that can be connected is determined by the number of ports available and other hardware constraints.

SANs enable expansion by scaling storage capacity across numerous systems and long distances. Scaling increases the number of devices and connections in a SAN. You can increase the number of switches in a fabric, or you can use routing technology to connect multiple SAN fabrics or multiple VSANs.

## Fabrics

A fabric is a single switch or a set of switches connected to form a network. Fabric services manage device names and addresses, timestamps, and other functionality for the switches.

A set of switches can be connected as a single fabric, an interconnected network of independent fabrics (LSANs for B-series), or partitioned into multiple logical fabrics (Virtual Fabrics for B-series or VSANs for C-series).

## SAN scaling

You can increase SAN connectivity by adding switches to an existing SAN or by using switches with more ports. When designing a SAN, you must ensure compliance with Fibre Channel standards and switch specifications. For switch-based scaling, consider the following factors:

- **Fibre Channel architecture**

Fibre Channel supports a maximum of 239 switches in a single fabric. HP specifies support based on rules for the maximum number of switches and maximum number of ports in a single fabric or multi-fabric SAN. Using many switches to obtain a high number of ports is unacceptable if the fabric exceeds the total switch count limit. Likewise, using large-capacity switches can create a network that exceeds the maximum number of ports.

For the HP-supported switch and port count fabric maximums, see “[B-series switches and fabric rules](#)” (page 83), “[C-series switches and fabric rules](#)” (page 115), “[H-series switches and fabric rules](#)” (page 131), and “[M-series switches and fabric rules](#)” (page 140).

- **Supported configurations**

Each Fibre Channel switch product line specifies the maximum number of ISLs, user ports, and hop counts, as well as link distances and other configuration limitations. The supported configurations determine the practical size of a SAN.

- **Fabric services**

Fabric services are distributed throughout the SAN to coordinate functions among all switches in the fabric. A large SAN requires the management functions provided by high-end switches. Some low-end switches have a limited capacity for expansion.

Routing technology facilitates SAN expansion beyond the capacity offered by switch-based scaling.

## Fibre Channel switches

A switch is identified by its function in a SAN:

- **Core** (or director)—Provides ISLs for any-to-any connectivity
- **Edge** (or fabric or SAN)—Provides user ports for connecting servers and storage systems

For some switches, the model name (for example, HP StorageWorks Core Switch 2/64) indicates its intended use in a SAN.

---

**NOTE:** This guide describes specific switch and fabric rules for SAN configuration. A heterogeneous environment requires coordination of components based on their rules to create a consolidated system. You must also consider the restrictions and requirements of the servers, HBAs, operating systems, cables, and other components.

---

## SAN design approaches

HP has three approaches to SAN design, listed here in order of complexity and experience required:

- **HP standard design**  
HP standard designs specify the arrangement of Fibre Channel switches in a SAN fabric, and are optimized for specific data access requirements and typical workloads. Implementing a standard design is the simplest approach to SAN design. HP recommends this approach for users who are designing a SAN for the first time.
- **Modified HP standard design**  
Select a standard SAN design that satisfies most of your requirements, and then modify it to meet your data access and connectivity requirements. HP recommends this approach for users with an intermediate level of SAN experience.
- **Custom design using the HP SAN design rules**  
Use a custom SAN design for specific storage and data access requirements. The SAN design rules in this guide specify guidelines for configuring custom topologies. HP recommends this approach for users with an intermediate or advanced level of SAN experience.

For information about:

- Standard SAN designs, see [“SAN fabric topologies”](#) (page 23)
- Customizing a SAN design, see:
  - [“B-series switches and fabric rules”](#) (page 83)
  - [“C-series switches and fabric rules”](#) (page 115)
  - [“H-series switches and fabric rules”](#) (page 131)
  - [“M-series switches and fabric rules”](#) (page 140)
- Heterogeneous SAN design, see:
  - [“Heterogeneous server rules”](#) (page 159)
  - [“P2000 and MSA storage system rules”](#) (page 203)
  - [“P6000/EVA storage system rules”](#) (page 213)
  - [“P9000/XP and VA storage system rules”](#) (page 227)
- Recommended SAN solutions and conventions, see [“Best practices”](#) (page 372)

## SAN design considerations

To design or modify a SAN, evaluate the following:

- **Geographic layout**  
The locations of campuses, buildings, servers, and storage systems determine the required SAN connections. SAN infrastructure components support long-distance connections and multiple interswitch cable segments. Fibre Channel routing interconnects independent SAN islands (fabrics) or VSANs to form a single, geographically distributed SAN.

For information about supported distances, see [“B-series switches and fabric rules”](#) (page 83).

- **Data availability**

A resilient SAN environment minimizes vulnerability to fabric or device failures and maximizes performance. A mixture of availability levels can be implemented in the same SAN, depending on the level of protection required for specific applications or data.

For information about availability levels, see [“Data availability”](#) (page 38).

- **Connectivity**

Provide enough ports to connect servers, storage systems, and fabric components. To create a high-capacity SAN, you can connect multiple fabrics or VSANs using routing.

For information about the connections available in a SAN fabric topology, see [“SAN fabric topologies”](#) (page 23).

- **Storage capacity**

Calculate the total storage capacity requirement and determine the type and number of storage systems needed for current and future requirements.

For storage systems information, see:

- [“P2000 and MSA storage system rules”](#) (page 203)
- [“P6000/EVA storage system rules”](#) (page 213)
- [“P9000/XP and VA storage system rules”](#) (page 227)

- **Heterogeneous platforms and operating systems**

Customize your SAN for specific hardware platforms and operating systems. In a heterogeneous environment, component interoperability depends on the capabilities and limitations of each platform.

For information about configuring systems in a heterogeneous environment, see [“Heterogeneous server rules”](#) (page 159).

- **Scalability and migration**

Choose a design that can be expanded incrementally over time as storage and connectivity needs increase. Migration paths for each of the topologies provide flexibility to expand a SAN. Fibre Channel routing accommodates expansion with minimal disruption to the network, especially where growth requirements are not known.

For information about scaling and migrating, see [“Best practices”](#) (page 372).

- **Backup and restore**

Provide adequate connectivity and bandwidth to maximize the performance of SAN-based backup.

For information about centralized backup, see [“Enterprise Backup Solution”](#) (page 247).

- **Disaster tolerance**

Consider remote data replication requirements to ensure protection against site failures and recovery of critical data.

For information about disaster tolerance and failover protection, see [“SAN extension”](#) (page 249).

- **Switch and hop counts**

Minimize the number of hops between devices that communicate regularly in the SAN.

For information about switches and hop counts, see:

- [“B-series switches and fabric rules” \(page 83\)](#)
- [“C-series switches and fabric rules” \(page 115\)](#)
- [“H-series switches and fabric rules” \(page 131\)](#)
- [“M-series switches and fabric rules” \(page 140\)](#)

- **Oversubscription**

For improved performance, reduce the potential for oversubscription. Ensure that the SAN design provides an adequate number of ISLs between switches, and minimize cases where many devices share a single-switch ISL.

For information about oversubscription, see [“Recommended ISL ratios” \(page 31\)](#).

- **Data locality, performance, and application workloads**

Provide an adequate level of performance based on application workloads. For frequent data reference and quick response times, use local, high-capacity paths to connect servers and storage systems. Deploy servers and storage in your SAN based on your data access requirements.

See [“SAN fabric topologies” \(page 23\)](#).

- **Manageability**

To enhance efficiency, you can manage consolidated storage from a centralized location.

- **Fabric zoning**

You can use fabric zoning to control SAN access at the device or port level.

For information about zoning, see [“B-series switches and fabric rules” \(page 83\)](#), [“C-series switches and fabric rules” \(page 115\)](#), [“H-series switches and fabric rules” \(page 131\)](#), and [“M-series switches and fabric rules” \(page 140\)](#).

- **Selective Storage Presentation**

To provide data access security and enable storage system use by multiple operating systems in a single SAN, use SSP.

- **SAN security**

Use a combination of SAN features and sound management practices to ensure data security throughout the SAN.

- **Fibre Channel routing functionality**

To increase the number of devices accessible in a SAN, use Fibre Channel routing functionality to interconnect existing SAN fabrics, Virtual Fabrics, or VSANs.

For routing functionality information, see [“SAN fabric topologies” \(page 23\)](#).

- **Virtual Fabrics (B-series switches) and virtual SANs (C-series switches)**

To create a SAN consisting of multiple logical SANs with separate fabric services, implement logical fabrics or VSANs. Use the IFR or inter-VSAN routing feature to enable device sharing across Virtual Fabrics or VSANs.

For information about Virtual Fabrics and VSANs, see [“SAN fabric topologies” \(page 23\)](#).

---

## 2 SAN fabric topologies

This chapter discusses HP standard SAN fabric topologies. It describes the following topics:

- “Fabric topologies” (page 23)
- “Single-switch fabric” (page 24)
- “Cascaded fabric” (page 25)
- “Meshed fabric” (page 26)
- “Ring fabric” (page 27)
- “Core-edge fabric” (page 29)
- “Topology data access” (page 32)
- “Topology maximums” (page 33)
- “Routed fabric topologies” (page 35)
- “FCoE fabric topologies” (page 38)
- “Data availability” (page 38)
- “Topology migration” (page 41)

### Fabric topologies

A SAN fabric topology defines the arrangement of Fibre Channel switches in a fabric. This section describes the HP-supported SAN fabric topologies.

There are three approaches to designing a SAN. You can implement:

- An HP standard SAN fabric topology design
- A subset or variation of an HP standard SAN fabric topology design
- A custom SAN fabric topology design

Regardless of which approach you use, the SAN design must adhere to the SAN design rules described in the following chapters:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “H-series switches and fabric rules” (page 131)
- “M-series switches and fabric rules” (page 140)
- “SAN fabric connectivity and switch interoperability rules” (page 146)
- “Heterogeneous server rules” (page 159)
- “P2000 and MSA storage system rules” (page 203)
- “P6000/EVA storage system rules” (page 213)
- “P9000/XP and VA storage system rules” (page 227)

## Routed SAN fabrics

HP standard fabric topologies support Fibre Channel routing. Fibre Channel routing enables connectivity between devices in multiple fabrics, Virtual Fabrics, or multiple VSANs. HP supports the following routed fabric technologies:

- “B-series Meta SAN” (page 35)
- “C-series VSANs with IVR” (page 37)
- “H-series switches with TR” (page 37)

## FCoE SAN fabrics

HP standard fabric topologies can integrate with FCoE technology. FCoE is deployed in existing Ethernet and Fibre Channel environments providing convergence at the server and fabric edge using CNAs and FCoE CN switches. See “FCoE fabric topologies” (page 38).

## Benefits


With HP standard SAN fabric topologies, you can:

- Create a SAN fabric for each department or application in your organization.
- Perform centralized management and backups.
- Update a SAN fabric to accommodate changing capacity or data access needs. You can also convert to another SAN fabric topology as needed.
- Connect devices over long distances using extended Fibre Channel or IP connections. See “SAN fabric connectivity rules” (page 146) and “SAN extension” (page 249).
- Connect multiple SAN fabrics using routing technology. See “B-series Meta SAN” (page 35).
- Deploy multiple logical fabrics using the Virtual Fabrics feature.
- Deploy multiple VSANs. See “C-series VSANs with IVR” (page 37).
- Incorporate a range of SAN availability levels. See “Data availability” (page 38).

## Single-switch fabric

A single-switch fabric consists of a Fibre Channel switch, server, and storage system (Figure 1 (page 24)). This topology forms the basis for all HP standard topologies. For example, you can connect two single-switch fabrics to create a cascaded fabric. Or, you can connect three or more single-switch fabrics to create a ring fabric or a core-edge fabric.

**Figure 1 Single-switch fabric**



## Switch models

For a small, single-switch SAN fabric, use the SN6000 Fibre Channel Switch; the 8/20q Fibre Channel Switch (or the HP Simple SAN Connectivity Kit); or an HP SAN, Fabric, or Edge switch (4, 8, 16, or 20 ports). For a larger single-switch SAN fabric, use a SAN, Fabric, or Edge switch



(32 to 40 ports), or a Core or Director switch (64 to 240 ports), which have higher port counts. For a high-availability SAN, use two switches configured in a dual-fabric SAN.

## Benefits

The benefits of a single-switch fabric include:

- Easy installation and configuration of servers and storage
- Maximum fabric performance because all communicating devices connect to the same switch
- Support for local, centralized, and distributed data access needs


## Cascaded fabric

A cascaded fabric is a set of interconnected switches, arranged in a tree format, that have one or more ISLs (Figure 2 (page 25)). You can connect one switch to one or more switches using a single ISL to each, or connect a pair of ISLs between two switches. HP recommends that you have a minimum of two ISL connections on each switch to provide fabric path redundancy. You should consider using a cascaded fabric topology if you require multiple groups of devices with localized intraswitch access.

Cascading enables you to:

- Achieve optimum I/O activity by connecting servers and storage to the same switch in the cascaded fabric
- Easily scale the fabric over time by adding cascaded switches

**Figure 2 Cascaded fabric**



## Switch models

All HP Fibre Channel switches are supported for use in a cascaded fabric topology. Cascaded fabric topologies typically use the SN6000 Fibre Channel Switch; the 8/20q Fibre Channel Switch (or the HP Simple SAN Connectivity Kit); or SAN, Fabric, or Edge switches, which support smaller incremental growth.

**NOTE:** Over time, a cascaded fabric topology can result in increased hops between switches. B-series, C-series, and H-series fabrics must not exceed seven hops; M-series fabrics must not exceed three hops. For additional switch hop information, see:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “H-series switches and fabric rules” (page 131)
- “M-series switches and fabric rules” (page 140)

## Benefits


The benefits of a cascaded fabric include:

- Ability to connect SANs in diverse geographic locations
- Ease of scalability for increased server and storage connectivity
- Shared backup and management support
- Optimum local performance when communicating devices are connected to the same switch in the cascaded fabric
- Cost efficiency due to the large number of switch ports available
- Support for local data access and occasional centralized data access

## Meshed fabric

A meshed fabric is a group of interconnected switches using multiple ISLs for fabric resiliency (Figure 3 (page 26)). If one ISL fails, the switch automatically reroutes data through an alternate path in the fabric. If the alternate path includes other switches, the data must pass through those switches to reach its destination.


**Figure 3 Meshed fabric**



25091a

As you add switches, ISLs are connected to two or more adjacent switches to maintain mesh connectivity, ensuring path redundancy throughout the fabric (Figure 4 (page 27)). The additional ISL connectivity provides communicating devices with more paths through the fabric. This dramatically reduces the chance that, as you add switches, you will exceed the maximum hop count.

**Figure 4 ISL connections in a meshed fabric**



## Switch models

All HP Fibre Channel switches are supported for use in a meshed fabric topology. Meshed fabric topologies typically use the SN6000 Fibre Channel Switch; the 8/20q Fibre Channel Switch (or the HP Simple SAN Connectivity Kit); or SAN, Fabric, or Edge switches, which support smaller incremental growth. To meet higher port-count requirements, use Core or Director switches.

## Benefits

The benefits of a meshed fabric include:

- Ability to meet multiple data access needs
- Multiple paths for internal fabric resiliency
- Ease of scalability
- Shared backup and management support
- Support for a mix of local and distributed data access (see “Topology data access” (page 32))
- Less impact on performance due to intraswitch traffic

## Ring fabric

A ring fabric is a ring of interconnected switches (Figure 5 (page 28)). The ring fabric provides a similar level of fabric resiliency as the meshed fabric and ensures full fabric connectivity with a minimum of two paths for each switch.

The ring fabric enables you to:


- Scale the fabric in a modular fashion.
- Achieve optimum I/O performance by connecting a group of servers and storage to one switch.

---

**NOTE:** HP does not recommend the ring fabric for applications requiring many-to-many connectivity.

---

**Figure 5 Ring fabric**



25093a

If the ring fabric has fewer than 12 switches, you can add switches (called satellite switches) outside the ring to create more user ports (Figure 6 (page 29)). M-series switches have a 3-hop maximum, reducing the total number of switches that can be configured in a ring to 7; satellite switches are not supported.

---


**NOTE:** Adding satellite switches slightly reduces fabric availability.

---

For more information on switch fabric maximums, see:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “H-series switches and fabric rules” (page 131)
- “M-series switches and fabric rules” (page 140)

**Figure 6 Ring fabric with satellite switches**



25094a

## Switch models

All HP Fibre Channel switches are supported for use in a ring fabric topology. Ring fabric topologies typically use the SN6000 Fibre Channel Switch; the 8/20q Fibre Channel Switch (or the HP Simple SAN Connectivity Kit); or SAN, Fabric, or Edge switches, which support smaller incremental growth. To meet higher port-count requirements, use Core or Director switches.

## Benefits

The benefits of a ring fabric include:

- Modular design and ease of scalability by adding a switch and other devices
- Multiple paths for internal fabric resiliency
- Support for a mix of local data access and occasional centralized data access

## Core-edge fabric


HP recommends using a core-edge fabric wherever possible.

A core-edge fabric has one or more Fibre Channel switches (called core switches) that connect to edge switches in the fabric (Figure 7 (page 30)). The core switches provide high bandwidth and redundant connectivity to the edge switches. The edge switches provide user ports for servers and storage. You can also connect centralized storage (disk or tape) to the core switches if centralized access is required.

The core-edge fabric is optimal for:


- Many-to-many connectivity environments that require high performance
- Unknown or changing I/O traffic patterns
- SAN-wide storage pooling

**Figure 7 Core-edge fabric (typical depiction)**



Core-edge fabric topologies are typically depicted as shown in [Figure 7 \(page 30\)](#), but can also be depicted hierarchically as shown in [Figure 8 \(page 30\)](#). Both figures represent the same physical implementation. How a topology is logically represented can help you understand the potential performance of a core-edge topology.

**Figure 8 Core-edge fabric (hierarchical depiction)**



## Core-edge fabric types

The number of ISLs between edge and core switches—typically expressed as a fan-in ratio, such as 7:1—characterizes the core-edge fabric types. The first number (7) indicates the number of edge ports. The second number (1) indicates the number of ISLs used by the edge ports to connect to a core switch in the fabric.

### Fat and skinny trees

There are two core-edge fabric topology types: fat tree and skinny tree. [Table 1 \(page 31\)](#) describes fat and skinny trees.

**Table 1 Core-edge fabric topology types**

Topology type	Description
Fat tree	At least 50% of edge ports are dedicated as ISLs, resulting in an ISL ratio of 1:1.
Skinny tree	Less than 50% of edge ports are dedicated as ISLs, resulting in an ISL ratio of x:1, where x is 2 or more.

### Recommended ISL ratios

The core-edge fabric type has a high fabric cross-sectional bandwidth (the maximum amount of data that can pass through ISLs at the fabric midpoint, which is the central connection or core of the fabric). The higher the ISL ratio, the lower the cross-sectional bandwidth and the more prone a topology is to ISL oversubscription. Oversubscription occurs when traffic is blocked due to insufficient ISL bandwidth.

**NOTE:** When determining the ideal ISL ratio, you must consider the speed of the server, storage, and ISL ports.

The minimum ISL ratio for an implementation depends on several factors, including:

- Location of server and storage fabric connection
- Server and storage hardware performance
- Data access type (see “Topology data access” (page 32))
- Server application performance requirements

Table 2 (page 31) describes the recommended core-edge fabric ISL ratios.

**Table 2 Recommended core-edge fabric ISL ratios**

I/O workload	Recommended ratios
Higher I/O data intensive application requirements (> 70 MB/s at 2 Gb/s, > 140 MB/s at 4 Gb/s, > 280 MB/s at 8 Gb/s)	1:1 to 3:1
Lower I/O data intensive application requirements (< 70 MB/s at 2 Gb/s, < 140 MB/s at 4 Gb/s, < 280 MB/s at 8 Gb/s)	7:1 to 15:1

**NOTE:** HP recommends a ratio of 7:1 for typical distributed data access.


### Numeric representation

Core-edge fabrics can also be represented in numeric terms, such as  $n1 \times n2$ , where  $n1$  represents the number of core switches and  $n2$  represents the number of edge switches.

For example, a 4 x 24 core-edge fabric indicates 4 core switches and 24 edge switches, for a total of 28 switches.

Figure 9 (page 32) shows a 4 x 12 core-edge fabric with 4 core switches and 12 edge switches, each connected to the core with 4 ISLs.

Figure 9 Core-edge fabric (4 x 12)



25097a

## Switch models

All HP Fibre Channel switches are supported for use in a core-edge fabric topology. Core-edge topologies typically use the SN6000 Fibre Channel Switch; the 8/20q Fibre Channel Switch (or the HP Simple SAN Connectivity Kit); SAN, Fabric, or Edge switches on the edge; and Core and Director switches in the core. H-series core-edge topologies use the SN6000 Fibre Channel Switches as edge and core switches (particularly if using the stacking capability) and 8/20q Fibre Channel Switches as edge and core switches. When using switches with different Fibre Channel maximum speed capabilities (such as 1 Gb/s, 2 Gb/s, 4 Gb/s, or 8 Gb/s), HP recommends using the higher-speed switches in the core.

## Benefits

The benefits of a core-edge fabric include:

- Typically, a maximum of two hops between switches
- Equal, centralized access to devices in the core
- Increased fabric and switch redundancy with two or more switches in the core
- Full many-to-many connectivity with evenly distributed bandwidth
- Support for centralized and distributed data access needs
- Ability to designate an optimally located core switch as the primary management switch, with direct connections to all switches

## Topology data access

To choose a SAN fabric topology, you must determine which data access type is appropriate for your environment. The data access types are as follows:

- **Local (one-to-one)**—Data access between a local server and a storage system connected to the same switch
- **Centralized (many-to-one)**—Data access between multiple, dispersed servers and one centrally located storage system
- **Distributed (many-to-many)**—Data access between multiple, dispersed servers and storage systems



Table 3 (page 33) lists the data access performance ratings for each SAN fabric topology.

**Table 3 Data access performance by SAN fabric topology**

SAN topology	Data access performance		
	Local	Centralized	Distributed
Single-switch fabric	Highest	Highest	Highest
Cascaded fabric	Highest	Not recommended	Not recommended
Meshed fabric	Medium	Medium	High
Ring fabric	Highest	Medium	Not recommended
Core-edge fabric (15:1, 7:1)	Medium	High	High
Core-edge fabric (3:1, 1:1)	High	Highest	Highest

## Topology maximums

Table 6 (page 34), Table 4 (page 33), Table 5 (page 34), and Table 7 (page 34) describe the maximum number of supported switches and ports for specific fabric topologies. In some cases, the number may be less than the maximums specified in the switch and fabric rules chapters. These differences relate to the number of hops in the fabric topology, as well as the number of ISLs, which affects the number of available user ports.

Consider the following:

- User ports are for server and storage connections.
- It is assumed that you have the minimum number of ISLs. If you require more ISLs, this reduces the number of user ports available for server and storage connections. See the following chapters for configuration limits:
  - “B-series switches and fabric rules” (page 83)
  - “C-series switches and fabric rules” (page 115)
  - “H-series switches and fabric rules” (page 131)
  - “M-series switches and fabric rules” (page 140)
- If you connect a Storage Management Appliance to the fabric, this further reduces the number of ports available for server and storage connections.

## B-series switches

Table 4 (page 33) lists the B-series switch and port maximums for specific fabric topologies.

**Table 4 B-series switch and port topology maximums**

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	512	512
Cascaded fabric	56	2,560	2,300
Meshed fabric			
Ring fabric	15		
Ring fabric with satellite switches	56		1,212
Core-edge fabric			2,300

## C-series switches

Table 5 (page 34) lists the C-series switch and port maximums for specific fabric topologies.

**Table 5 C-series switch and port topology maximums**

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	528	528
Cascaded fabric	60	4,000 (maximum of 12 Director switches)	3,500 (cascaded with 12 Director switches and 10 Fabric switches)
Meshed fabric	60		3,500
Ring fabric	15		
Ring fabric with satellite switches	60		
Core-edge fabric	60		

## H-series switches

Table 6 (page 34) lists the H-series switch and port maximums for specific topologies.

For large fabrics, the number of user ports is determined by use of SN6000 Fibre Channel Switches and making maximum use of their dedicated 10Gb/20Gb stacking port ISLs, which leaves more of the 8 Gb ports available for device connections.

**Table 6 H-series switch and port topology maximums**

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	20	20
Cascaded fabric	30	600	460
Meshed fabric			512
Ring fabric	15	300	264
Ring fabric with satellite switches	30	600	462
Core-edge fabric			

## M-series switches

Table 7 (page 34) lists the M-series switch and port maximums for specific fabric topologies.

**Table 7 M-series switch and port topology maximums**

SAN topology	Number of switches	Total number of ports	Number of user ports
Single-switch fabric	1	140	140
Cascaded fabric	24 (maximum of 8 Director switches)	1,632	1,024 (cascaded with 8 Director switches and 16 Edge switches)
Meshed fabric	N/A	N/A	N/A
Ring fabric	7	980	966

**Table 7 M-series switch and port topology maximums** *(continued)*

SAN topology	Number of switches	Total number of ports	Number of user ports
Ring fabric with satellite switches	N/A	N/A	N/A
Core-edge fabric	24	1,632	1,024

## Routed fabric topologies

HP standard fabric topologies support Fibre Channel routing, which provides connectivity between devices in multiple fabrics or VSANs.

This section describes the following HP Fibre Channel routed fabric technologies:

- [“B-series Meta SAN” \(page 35\)](#)—Implemented in certain 8 Gb/s switch models with license enabled integrated Fibre Channel routing or using the B-series 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 Multi-protocol Router (400 MP Router), Multi-protocol Router Blade (MP Router Blade), or Multi-protocol Router (MP Router), which provide selective Fibre Channel routing connectivity between multiple B-series fabrics or between B-series and M-series fabrics.
- [“B-series Virtual Fabrics with IFR” \(page 36\)](#)—Implemented in B-series switches (DC SAN Backbone Director, DC04 Director, 8/80 SAN Switch, and 8/40 SAN Switch) using B-series Virtual Fabrics. IFR allows you to share devices across multiple B-series Virtual Fabrics partitions.
- [“C-series VSANs with IVR” \(page 37\)](#)—Implemented using C-series IVR. IVR provides selective Fibre Channel routing connectivity between devices in different VSANs. Ports on one or more switches can be assigned to different VSANs.
- [“H-series switches with TR” \(page 37\)](#)—Implemented using the TR feature, which is available with firmware 8.x (or later). The TR feature provides inter-fabric routing, allowing controlled access between devices on an SN6000 Fibre Channel Switch or an 8/20q Fibre Channel Switch (local) fabric and devices on a remote fabric consisting of B-series or C-series switches.

## B-series Meta SAN

A Meta SAN contains multiple B-series fabrics connected together using the B-series 8 Gb/s switches with Fibre Channel routing, 1606 Extension SAN Switches, DC Director Switch MP Extension Blades, or MP Routers. The 8 Gb/s switches with Fibre Channel routing, 1606 Extension SAN Switches, DC Dir Switch MP Extension Blades, or MP Routers implement the Fibre Channel routing service, which allows selective access between devices in different fabrics without having to merge fabrics. This provides a high level of isolation between fabrics. This isolation can be viewed as individual Fibre Channel subnetworks within the Meta SAN.

LSANs provide access to devices in different fabrics. You create LSAN zones just as you create standard zones in a single fabric. The difference is that LSAN zone definitions span multiple fabrics and therefore must be replicated on all fabrics that comprise the LSAN.

For more information about Meta SANs and Fibre Channel routing, see [“Fibre Channel routing” \(page 43\)](#).

## B-series Virtual Fabrics architecture

In B-series configurations, a switched fabric is partitioned into several logical switches and logical fabrics ([Figure 15 \(page 47\)](#)) by using the Virtual Fabrics feature. To accomplish this, the Virtual Fabrics feature must be enabled. The Virtual Fabrics feature is set to `OFF` by default. Enabling Virtual Fabrics is a disruptive operation, which requires a reboot. Devices can be shared across multiple logical fabrics using IFR, increasing resource sharing. B-series Virtual Fabrics and Virtual Fabrics with IFR are included on B-series switch models: DC SAN Backbone Director, DC04 Director, 8/80 SAN Switch and 8/40 SAN Switch only.

## B-series fabric partitioning with administrative domains

In B-series configurations, a switched fabric is partitioned into several management domains using the Administrative Domains feature. An administrative domain is only a filtered administrative view of the fabric. Each administrative domain has its own set of resources, such as: administrator and users, Name Server, and zoning database. Resources can be shared across multiple administrative domains by properly defining their membership. The Administrative Domains feature is available on all switches using firmware 5.2x (or later).

---

**NOTE:** Virtual Fabrics and Administrative Domains are mutually exclusive; they cannot run on the switch simultaneously. To use Administrative Domains, you must first disable Virtual Fabrics; to use Virtual Fabrics, you must first delete all Administrative Domains.

---

## Switch models and fabric topologies

HP supports Meta SANs with all B-series switches in either HP standard or customized topologies. You must follow all B-series and Meta SAN fabric rules. For more information about Meta SANs, see [“B-series switches and fabric rules” \(page 83\)](#).

### Benefits

A Meta SAN:

- Allows fabric connections (without the need to merge fabrics), providing a high level of fault isolation and centralized fabric management
- Connects multiple SAN islands (independent fabrics), enabling selective resource sharing
- Eliminates the need to move and re-cable equipment in different fabrics
- Allows connection of fabrics with the same domain ID and zoning definitions
- Reduces the impact of scaling limits for individual fabrics
- Increases levels of storage consolidation
- Provides centralized backup for multiple fabrics
- Allows higher level of fabric management and management consolidation

## B-series Virtual Fabrics with IFR

The B-series Virtual Fabrics feature allows you to partition fabrics. You create partitions in the fabric by creating logical switches within a physical switch and by creating logical fabrics within a physical switch or across multiple switches. This provides a high level of isolation between Virtual Fabrics partitions or logical fabrics, allowing you to view individual Fibre Channel subnetworks in a B-series fabric.

The IFR feature allows you to configure devices in one logical fabric for access to devices in another logical fabric. B-series Virtual Fabrics and Virtual Fabrics with IFR are included on selected 8 Gb switches only.

## Switch models and fabric topologies

HP supports Virtual Fabrics in either HP standard or customized topologies. You must follow all B-series fabric rules. For more information, see [“B-series switches and fabric rules” \(page 83\)](#).

### Benefits

Virtual Fabrics:

- Isolate fabric services and minimize fault propagation
- Allow multiple secure fabric partitions over the same physical infrastructure

- Restrict device access for improved control and security
- Provide selective device access and sharing using the IFR feature

## C-series VSANs with IVR

VSANs are groups of switch ports from one or more C-series switches. Each VSAN has a unique set of fabric services. Different fabric settings can be applied to each VSAN. This provides a high level of isolation between VSANs. This isolation can be viewed as individual Fibre Channel subnetworks within a C-series fabric.

The IVR feature enables you to configure devices in one VSAN for access to devices in another VSAN. All C-series switches include the VSAN feature. IVR is an optional licensed software feature. There is no need for additional hardware.

### Switch models and fabric topologies

HP supports VSANs with all C-series switches in either HP standard or customized topologies. You must follow all C-series and VSAN fabric rules.

For more information about VSANs, see [“C-series switches and fabric rules”](#) (page 115).

### Benefits

A VSAN:

- Isolates fabric services and minimizes fault propagation
- Allows multiple secure VSANs over the same physical infrastructure
- Restricts device access for improved control and security
- Provides selective device access and sharing using the IVR feature

## H-series switches with TR

The TR feature provides inter-fabric routing on a per-port basis, allowing controlled access between devices on an H-series switch (local) fabric and devices on a remote fabric consisting of B-series or C-series switches. The establishment of a routed connection using TR maintains a high level of isolation between fabrics. A transparent route between two devices consists of a connection from a TR\_Port on an SN6000 Fibre Channel Switch or an 8/20q Fibre Channel Switch to a switch in the remote fabric, a mapping of the two devices to be routed together, and an IFZ for the routed devices in both fabrics.

A TR\_Port server as a bridge between the transparent router's local fabric and a remote fabric. The TR\_Port uses standard NPIV login methods to connect to the remote fabric. You can configure any of the 8 Gb ports of the SN6000 Fibre Channel Switch or the 8/20q Fibre Channel Switch as TR\_Ports.

### Switch models and fabric topologies

HP supports the SN6000 Fibre Channel Switch or the 8/20q Fibre Channel Switch with TR with B-series or C-series switches in either HP standard or customized topologies. For the supported B-series and C-series switches and firmware, see [“Supported switches in an H-series switch with TR remote fabric”](#) (page 137). You must follow all H-series switch fabric rules. For more information, see [“H-series switches and fabric rules”](#) (page 131).

### Benefits

H-series switches with TR provide the following benefits:

- Allows fabric connections (without the need to merge fabrics), providing a high level of fault isolation and centralized fabric management
- Allows connection of fabrics with the same domain ID and zoning definitions

- Reduces the impact of scaling limits for individual fabrics
- Increases levels of storage consolidation
- Provides centralized backup for multiple fabrics

## FCoE fabric topologies

FCoE technology allows you to converge Ethernet and Fibre Channel technology, providing significant cable, adapter, and switch consolidation. All HP-supported Fibre Channel topologies are supported integrated with FCoE at the Fibre Channel fabric edge. HP also supports end-to-end FCoE solutions for the P6300/P6500 EVA storage systems with the FCoE target interface and P6000 (formerly EVA) storage using the MPX200 Multifunction Router. For more information about FCoE, see [“Fibre Channel over Ethernet” \(page 61\)](#).

## Data availability

SAN data availability depends on the reliability of the SAN fabric, servers, and storage systems during routine operations. The data availability level required for your SAN environment is based on:

- Administrative requirements (for example, backup schedules, operating procedures, and staffing)
- Protection level for applications or data
- Hardware redundancy

---

**NOTE:** For more information about high-availability configurations when using the HP B-series MP Routers or C-series VSANs, see [“High-availability router configurations” \(page 55\)](#).

---

Several factors affect SAN data availability:

- Application software
- Server operating systems
- Server hardware
- SAN fabric infrastructure
- Primary and secondary storage
- Number of switches
- Number of ISLs
- Number of paths between a server or clustered servers and the fabric
- Number of storage controller paths in the fabric


## Levels

This section describes the data availability levels.

### Level 1: single connectivity fabric

Level 1 provides maximum connectivity but does not provide fabric resiliency or redundancy. Each switch has one path to other switches in the fabric ([Figure 10 \(page 39\)](#)). Each server and storage system has one path to the fabric.


**Figure 10 Level 1: single connectivity fabric**



### Level 2: single resilient fabric

Level 2 provides fabric path redundancy by using multiple ISLs between switches (Figure 11 (page 39)). Each server and storage system has one path to the fabric. If an ISL or switch port failure occurs, the switch automatically reroutes data through an alternate fabric path and there is no interruption in server I/O activity.

**Figure 11 Level 2: single resilient fabric**



### Level 3: single resilient fabric with multiple device paths

Level 3 is the same as level 2 but also provides multiple server and storage system paths to the fabric to increase availability (Figure 12 (page 40)). If a switch, server HBA, or storage system path failure occurs, data is automatically rerouted through an alternate path and there is no interruption in server I/O activity.


To take full advantage of this level, HP recommends that you connect each server HBA and each storage system path to a different switch to increase availability and reduce the potential for an SPOF. This level provides both fabric resiliency and device path redundancy.

---

**NOTE:** Certain operating systems may require the use of fabric zoning to define a minimum of two zoned paths for each server configured with multiple paths in a single fabric.

---

**Figure 12 Level 3: single resilient fabric with multiple device paths**

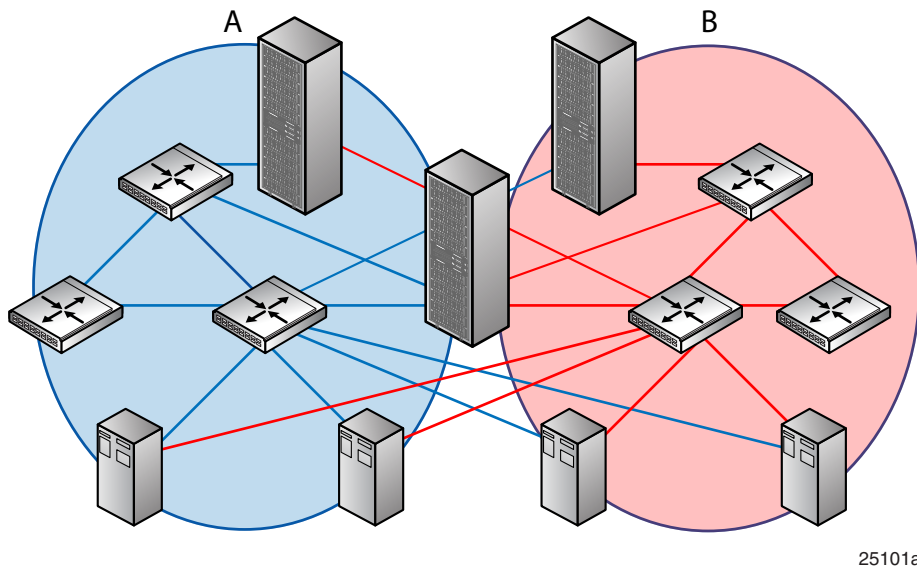


#### Level 4: multiple fabrics and device paths (NSPOF)

Level 4 provides multiple data paths between servers and storage systems, but unlike level 3, the paths connect to physically separate fabrics (Figure 13 (page 40)). This level ensures the highest availability with NSPOF protection. If a switch, server HBA, or storage system path failure occurs, data is automatically rerouted through the alternate fabric and there is no interruption in server I/O activity.

Level 4 minimizes vulnerability to fabric failures (for example, improper switch replacement, incorrect fabric configuration settings, or a fabric service failure). Level 4 also provides the highest level of performance and a higher number of available ports, since all fabrics can be accessed simultaneously during normal operations.

**Figure 13 Level 4: multiple fabrics and device paths (NSPOF)**



Using two fabrics may increase implementation costs, but it also increases the total number of available ports. For example, in a single meshed fabric with four switches, you have a maximum of 52 user ports for servers and storage. Implementing the same topology using two fabrics increases the maximum number of user ports to 104.



## Considerations

When choosing a data availability level, you must consider:

- Cost
- Access to critical data

For mission-critical applications, HP recommends that you implement a level 4, fully redundant fabric configuration. You can justify the additional cost if you consider the cost of losing access to critical data.

Table 8 (page 41) indicates data availability and supported topologies for each level.

**Table 8 Fabric design data availability**

Fabric design	Availability level	SAN topologies
Level 1: single connectivity fabric	No redundancy	Single switch or multiple switches with single ISL
Level 2: single resilient fabric	Medium	Cascaded with two ISLs, meshed, ring, and core-edge
Level 3: single resilient fabric with multiple device paths	High	All
Level 4: multiple fabrics and device paths (NSPOF)	Highest	All

You can add fabrics to increase the number of available ports. Table 9 (page 41) lists the cost calculations for each data availability level.

**Table 9 Calculating for data availability levels**

Fabric design	Hardware cost	Number of available ports
Level 1: single connectivity fabric	$x^1$	# ports = $n$ – number of ISL ports <sup>2</sup>
Level 2: single resilient fabric	$x$ + additional ISLs	# ports = $n$ – number of ISL ports
Level 3: single resilient fabric with multiple device paths	$x$ + additional ISLs + additional HBAs	# ports = $n$ – number of ISL ports – additional HBA ports <sup>3</sup>
Level 4: multiple fabrics and device paths (NSPOF)	$x$ + additional ISLs + additional HBAs + additional switches	# ports = $2n$ – number of ISL ports – additional HBA ports

<sup>1</sup>  $x$  is the cost of a single connectivity fabric.

<sup>2</sup>  $n$  is the total number of ports for servers and storage systems.

<sup>3</sup> May require use of zoning to define a minimum of two data paths in a single fabric (operating system dependent).

## Topology migration

To increase SAN connectivity and capacity:

- Increase the number of switches.
- Use switches with more ports.
- Implement multiple fabrics.
- Implement Fibre Channel routing.
- Migrate to another fabric topology.
- Deploy multiple independent SANs.

## Nondisruptive migration

If you have a level 4 multiple fabric NSPOF SAN, you can fail over all operations to one fabric and then reconfigure the other fabric.

When planning a migration, try to avoid or minimize the movement of devices between switches. Migrations that require the addition or re-cabling of ISLs are less disruptive than migrations that require movement of device connections.

## Migrating a cascaded fabric SAN

This section describes migration paths for a cascaded fabric SAN.

### Cascaded to meshed

To create a meshed fabric SAN, you need additional ISLs to connect all switches. To ensure a successful migration, calculate the number of ports needed for the additional ISLs. You may need to move device connections to another switch to make ports available for ISLs.

### Cascaded to ring

If you have a linear cascaded fabric SAN, connect the last switch in the fabric to the first switch to create a ring fabric SAN. If you have a tree-like cascaded fabric SAN (with multiple levels), you may need to re-cable the ISLs.

### Cascaded to core-edge

Determine which switches will be the backbone switches and which ones will be the edge switches. Re-cable the ISLs to connect all edge switches to the core switches. Connect devices (servers and storage) or core switches, as required. This migration is less disruptive if you use the existing switches as edge switches and add switches as core switches.

## Migrating a meshed fabric SAN

This section describes migration paths for a meshed fabric SAN.

### Meshed to ring

You can migrate a meshed fabric SAN to a ring fabric SAN by removing the cross-connected ISLs and leaving the outer-connected ISLs as a ring. Use the available ports for device connections or for redundant ring ISL connections.

### Meshed to core-edge

Use the method described in [“Cascaded to core-edge”](#) (page 42).

## Migrating a ring fabric SAN

This section describes migration paths for a ring fabric SAN.

### Ring to meshed

If you have two ISLs between all switches in the ring fabric, re-cable each ISL so that it connects to the appropriate switch in the meshed fabric you design.

### Ring to core-edge

This migration is less disruptive if you have two ISLs between all switches in the ring fabric SAN. Use the method described in [“Cascaded to core-edge”](#) (page 42).

---

## 3 Fibre Channel routing

This chapter describes Fibre Channel routing in an HP SAN environment. It describes the following topics:

- “Fibre Channel routing overview” (page 43)
- “SAN scaling and routing” (page 44)
- “Fibre Channel routing implementations” (page 46)
- “Fabric redundancy and routing” (page 51)
- “Supported routing configurations” (page 52)

### Fibre Channel routing overview

Fibre Channel routing facilitates the development and management of higher-capacity SANs, significantly increasing device connectivity. By enabling communication between two or more physically independent fabrics, multiple logical fabrics, or VSANs, routing provides high levels of SAN scalability. Each fabric, logical fabric, or VSAN maintains a unique fabric services configuration.

---

**NOTE:** In the context of Fibre Channel routing, the terms "fabric," "Virtual Fabric," and "VSAN" are used interchangeably. HP does not support using the B-series Fibre Channel routing, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, Multi-protocol Routers, or B-series Virtual Fabrics-IFR with C-series IVR functionality in the same SAN configuration.

HP only supports routing devices connected to H-series switches through TR\_Ports to devices in B-series and C-series fabrics.

---

Routing enables independent fabrics, Virtual Fabrics with IFR, or VSANs with IVR to dynamically share devices without the need to reconfigure or re-cable physical connections.

Routed fabrics, Virtual Fabrics with IFR, or VSANs with IVR can consolidate management interfaces. Instead of one management interface per fabric, there can be one per SAN, or two per SAN, if redundant fabrics are used.

Routing using the TR feature of the H-series switches does not provide consolidated management of the routed fabric.

Fibre Channel routing features include:

- **Increased SAN scalability**
  - Interconnecting (not merging) multiple physical fabrics, Virtual Fabrics, or VSANs
  - Overcoming individual fabric scaling limits
- **Improved device access and sharing**
  - Sharing devices dynamically across multiple fabrics, Virtual Fabrics, or VSANs
  - Increasing device utilization
- **Fabric, Virtual Fabric, or VSAN independence**
  - Isolation of fault domains
  - Separate fabric services
- **Centralized SAN fabric management**
  - Common fabric management
  - Tape backup consolidation

## Fabric, Virtual Fabric, and VSAN independence

Fibre Channel routing identifies data frames in a fabric, Virtual Fabric, or VSAN for transfer to other fabrics, Virtual Fabrics with IFR, or VSANs with IVR. Only data addressed to a device in another fabric, Virtual Fabric, or VSAN passes through the router or routing function; therefore, a disruption of fabric services in one routed fabric, Virtual Fabric, or VSAN is unlikely to propagate to another.

## Fabric services

Fabric services coordinate communication between switches in a fabric, Virtual Fabric, or VSAN.

The fabric services manage:

- Device names and addresses
- Timestamps
- Switch utilities

Routing connects devices in multiple fabrics, Virtual Fabrics, or VSANs without extending fabric services from one routed fabric to another. Devices in a routed network can communicate across LSANs, Virtual Fabrics, or VSANs despite having different fabric services configurations.

## Worldwide Name

A recognized naming authority assigns each Fibre Channel device a unique identifier, called the WWN. Use the device WWNs to:

- Assign devices to zones.
- Define devices to export from one fabric, Virtual Fabric, or VSAN to another.

## Import and export

Routing creates a Meta SAN, extended Virtual Fabric, or extended VSAN when it connects fabrics, Virtual Fabrics, or VSANs. Routing exports devices from one fabric, Virtual Fabric, or VSAN to another. An exported device has an imported address in every destination fabric, Virtual Fabric, or VSAN to which it has been exported. The address of the exported device in the source fabric, Virtual Fabric, or VSAN is its exported address.

An imported device is a device as seen in a fabric when using its imported address. An exported device is a device as seen in the fabric when using its exported address.

## Routing table

The routing function reads the fabric address information in each frame that it receives, and then uses a routing table to determine the destination fabric, destination Virtual Fabric, or destination VSAN and the address within that fabric, Virtual Fabric, or VSAN. The routing function then transmits the frame to the address in the destination fabric.

## SAN scaling and routing

This section describes two methods for increasing the size of SANs:

- Increase the Fibre Channel switch capability within a fabric.
- Connect independent fabrics using a Fibre Channel router, Virtual Fabrics with IFR, or VSANs with IVR.

## Switch scaling

The switches that make up fabrics define the fabric limits. This section describes the relationship between switches.

## Switch scaling limits

Adding ports to a fabric means increasing the number of switches in the fabric or increasing the number of ports per switch. For large fabrics, adding ports may not be possible unless the limits for total port count and total switch count are increased.

Each Fibre Channel switch product line has its own limits for total port count and switch count. You must ensure that a new or modified SAN design complies with these limits.

---

**NOTE:** Other limits, such as hop counts and link distances, also apply. For more information, see:

- [“B-series switches and fabric rules” \(page 83\)](#)
  - [“C-series switches and fabric rules” \(page 115\)](#)
  - [“H-series switches and fabric rules” \(page 131\)](#)
  - [“M-series switches and fabric rules” \(page 140\)](#)
- 

For a SAN design to meet the total port count and total switch count limits, the following configuration restrictions are enforced:

- The fabric size limit for total port or total switch count must not be exceeded.
- The use of several small switches to reach a high total port count number is not acceptable if the design exceeds the total switch count limit.
- The use of several high-port-count switches is not acceptable if the design exceeds the total port count limit.

For fabric configurations, HP defines the maximum supported port and switch counts.

## Fabric services limits

Fabric services provide coordination between all switches in a fabric. Increasing fabric size increases the overhead associated with coordination. Fabric services include:

- Fabric Login Server
- State Change
- Notification Server
- Name/Directory Server
- Zone Server
- Key Server
- Time Server
- Simple Name Service

## Simple fabric services

SNS provides a mapping between device names and their addresses in a fabric. To ensure that the mapping is up-to-date, every switch in the fabric implements SNS.

## Coordinating fabric services

Each fabric maintains a unique set of fabric services. When two fabrics are connected, their two sets of services merge to form a single set.

As fabrics grow, coordinating the fabric services across switches, hosts, and storage devices becomes more challenging. It is difficult to match the fabric service requirements for very small, inexpensive switches with those for large, high-end switches. Without routing, fabric scaling is limited by the ability of the smallest fabric switch to participate in the distributed fabric services system.

## Scaling by routing

Increasing fabric port count and switch count limits meets most customer scaling requirements. Demand for higher port counts and connectivity between devices in different fabrics, Virtual Fabrics, or VSANs requires Fibre Channel routing.

Routing improves scaling by connecting independent fabrics, Virtual Fabrics, or VSANs, each potentially at its full capacity. Connectivity between fabrics, Virtual Fabrics, or VSANs allows sharing of resources, reducing unnecessary redundancy in the routed network.

You can route between fabrics without affecting the total switch and port count limits. However, the routed network is not the same as a single large fabric, Virtual Fabric, or VSAN. Only selected devices in each fabric, specified by a routing table, can communicate with devices in other fabrics.

For example, using a router, you can connect three 1,200-port fabrics to construct a 3,600-port Meta SAN. You determine which fabrics require connectivity, and then specify the devices allowed to communicate across fabrics. The router does not provide 100% any-to-any connectivity between fabrics, but it does meet most SAN requirements.

## Fibre Channel routing implementations

With Fibre Channel routing, you can create a routed fabric by:


- Connecting several fabrics using a router or a switch with router functionality
- Partitioning a fabric into several Virtual Fabrics
- Dividing a single fabric into several smaller Virtual Fabrics or VSANs

## Fibre Channel routing techniques

This section describes the following Fibre Channel routing techniques:


- B-series routing connects independent fabrics (SAN islands), as shown in [Figure 14 \(page 46\)](#). HP-supported B-series router products include:
  - 8 Gb/s switch models with license enabled integrated Fibre Channel routing
  - HP StorageWorks 1606 Extension SAN Switch (1606 Extension SAN Switch)
  - HP StorageWorks DC SAN Director Multi-protocol Extension Blade (DC Dir Switch MP Extension Blade)
  - 400 Multi-protocol Router (400 MP Router)
  - Multi-protocol Router Blade (MP Router Blade for the 4/256 SAN Director)
  - Multi-protocol Router (MP Router)
- A B-series switch with IFR connects multiple Virtual Fabrics, as shown in [Figure 15 \(page 47\)](#).
- A C-series switch with IVR connects multiple VSANs, as shown in [Figure 16 \(page 47\)](#).
- An H-series switch with TR connects to other B-series or C-series fabrics, as shown in [Figure 17 \(page 47\)](#).

**Figure 14 Basic MP Router configuration**




25102c


**Figure 15 Basic Virtual Fabric IFR configuration**



**Figure 16 Basic VSAN IFR configuration**



**Figure 17 Basic TR configuration**



## B-series fabric groups

In B-series routing configurations, devices in different fabrics can be grouped to form LSANs. An LSAN is similar to a Fibre Channel zone, but can extend through a router to include devices in other fabrics. This configuration, which includes the physical fabrics (subnetworks), LSANs, and router, is called a *Meta SAN*. A Meta SAN consolidates multiple fabrics into a single entity.

Figure 14 (page 46) shows Fabric 1, Fabric 2, and Fabric 3, each containing one or more switches. Any B-series switch can be used in these fabrics. In each fabric, the switches must use the same version of switch firmware for like switches and must have the same variable settings (for example, R\_A\_TOV). Each fabric has a unique set of fabric services. For fabric restrictions, see “B-series switches and fabric rules” (page 83).

Fabrics connected with routing must comply with configuration rules for a routed fabric. See “1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules” (page 100), “Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules” (page 103), and “MP Router fabric rules” (page 107). The fabrics can have identical domain names and zoning definitions.

The 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades and MP Routers also provide FCIP capabilities, allowing implementation of Fibre Channel routing and FCIP SAN extension. See “Integration of Fibre Channel routing and FCIP” (page 57).

## B-series fabric partitioning using Virtual Fabrics

In B-series configurations, a switched fabric is partitioned into several Virtual Fabrics (Figure 15 (page 47)). Each Virtual Fabric has its own set of resources, such as administrator and users, Name Server, and zoning database. Devices can be shared across multiple Virtual Fabric administrative domains using IFR, thus increasing resource sharing. B-series Virtual Fabrics and Virtual Fabrics with IFR are included on all B-series switches using firmware 5.2x (or later) without the need for a router.

## B-series Virtual Fabrics architecture

In B-series configurations, a switched fabric is partitioned into several logical switches and logical fabrics by using the Virtual Fabrics feature. See Figure 15 (page 47). Logical switches within a physical switch can be created by dividing the switch ports and assigning them to individual logical switches. An FID also must be configured to each logical switch. A logical fabric is a fabric that contains at least one logical switch; but logical switches can be connected to other logical switches with the same FID to form logical fabrics across multiple switches. Devices can be shared across multiple logical fabrics using IFR, thus increasing resource sharing. B-series Virtual Fabrics and Virtual Fabrics with IFR are included on B-series switch models: DC SAN Backbone Director, DC04 Director, 8/80 SAN Switch and 8/40 SAN Switch only.

## C-series fabric partitioning

In C-series configurations, a single fabric is partitioned into several subnetworks or logical groups of switches or switch ports called VSANs. The group of VSANs is called a SAN.

Figure 16 (page 47) shows VSAN 1, VSAN 2, and VSAN 3, each a set of switch ports on one or more C-series switches. A VSAN can extend across multiple switches. Each VSAN has a unique set of fabric services with independent fabric management. VSANs can share devices by using the license-enabled IVR function. IVR is distributed across all switches in the SAN, and there is no separate router hardware. Because the switches are a connected set, they must run the same version of switch firmware.

## H-series switch fabric routing

You can configure any H-series switch 8 Gb port as a TR\_Port, which you use to connect devices on the H-series switch to devices on a remote fabric. You do this by configuring TR mapping, which establishes a route to connect one device on the H-series switch to one device on a remote fabric through one TR\_Port. Multiple devices can share TR\_Ports, and you can configure multiple TR\_Ports to the same remote fabric. HP currently supports connection to B-series and C-series remote fabrics.

Figure 17 (page 47) shows how one or more remote fabrics can connect to an H-series switch. Remote Fabric 1, Remote Fabric 2, and the H-series switch fabric each contain one or more switches. Devices connected through routing must comply with the configuration rules for the TR function. See “Fabric rules for H-series switches with TR” (page 136). The fabrics can have identical domain names and zoning definitions.

## B-series, C-series, and H-series routing differences

B-series 8 Gb/s switches with integrated Fibre Channel routing, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, an MP Router, Virtual Fabrics with IFR, or VSANs with IVR can connect existing fabrics or VSANs. When existing fabrics are connected to an 8 Gb/s switch with Fibre Channel routing or MP Router, it creates a Meta SAN. Using B-series switches with Virtual Fabrics or C-series switches with VSANs, existing fabrics are physically connected, and the routing function in the switches is configured using IFR or IVR.

As shown in Figure 18 (page 49), an LSAN can include devices connected to different fabrics (for example, the LSAN Zone connects devices from Fabric 1 and Fabric 2).



**Figure 18 B-series routing**

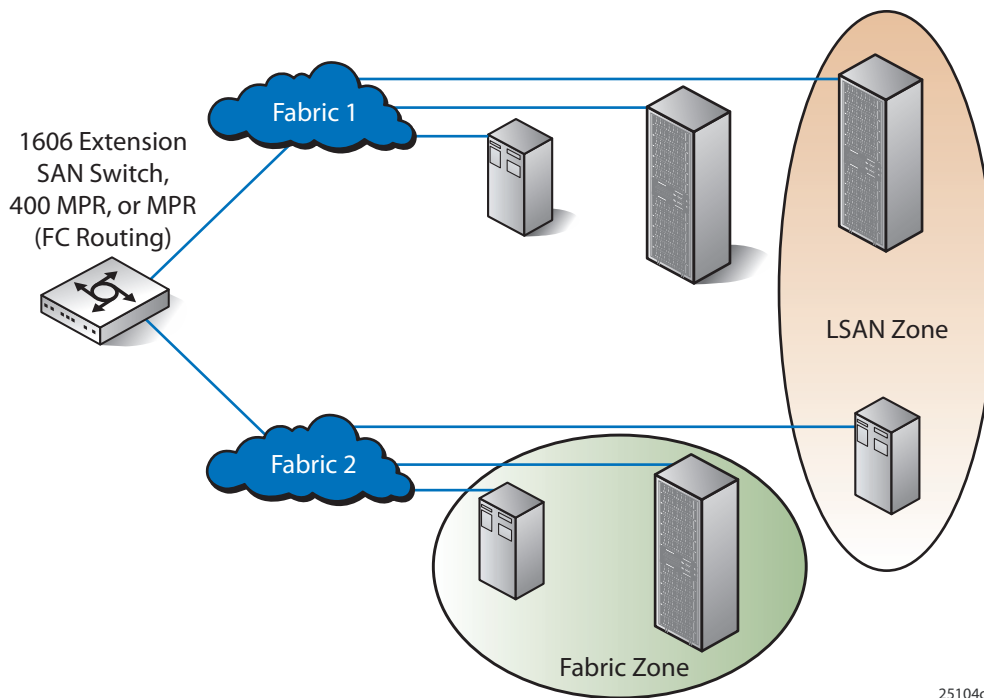



Figure 19 (page 49) and Figure 20 (page 50) show the differences between B-series MP Router and C-series routing.


Figure 19 (page 49) shows how Virtual Fabrics can include devices that connect to a single switch or multiple switches in the SAN. Devices in different Virtual Fabrics communicate using IFR. Multiple switches can be connected in any supported fabric configuration.

**Figure 19 B-series Virtual Fabric IFR**




As shown in [Figure 20 \(page 50\)](#), VSANs can include devices that connect to a single switch or multiple switches in the SAN. Devices in different VSANs communicate using IVR. Multiple switches can be connected in any supported fabric configuration.

**Figure 20 C-series VSAN IVR**



[Figure 21 \(page 50\)](#) shows how TR\_Ports on the H-series switch can connect devices in a local fabric to devices in a remote fabric. When a device on the H-series switch is mapped to a device in the remote fabric, the H-series switch automatically creates an inter-fabric zone whose members are the two devices and the TR\_Port that connects them. CLI commands for adding this zone to the remote fabric zone set are generated automatically.

**Figure 21 H-series switch routing**




## Fabric redundancy and routing

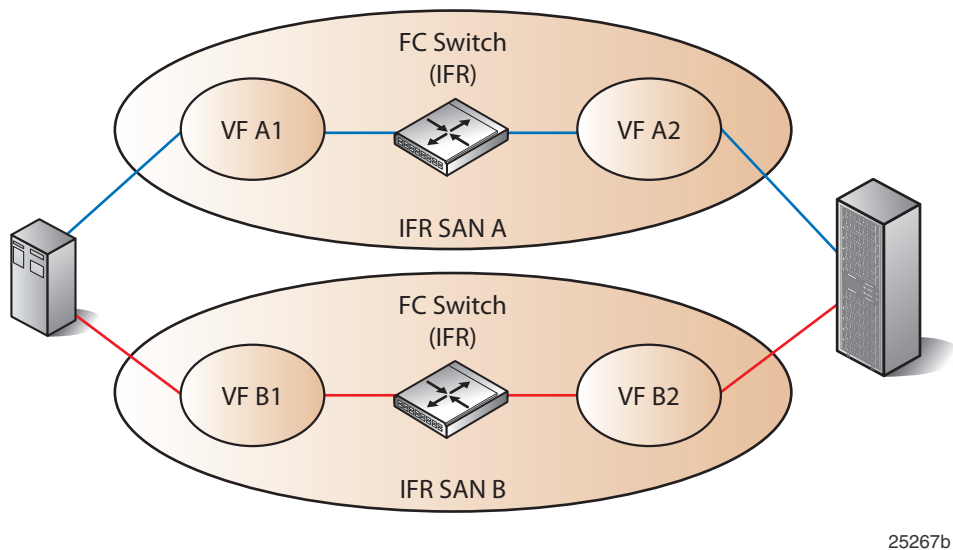
B-series 8 Gb/s switches with integrated Fibre Channel routing, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Routers can connect one group of fabrics, or a single fabric can connect multiple Virtual Fabrics or VSANs with IFR or IVR. For a high-availability, fully redundant implementation, you can have two routers and two groups of fabrics (Figure 22 (page 51)), or two groups of Virtual Fabrics or VSANs with IFR (Figure 23 (page 51)) or IVR (Figure 24 (page 52)).

For more information, see “High-availability router configurations” (page 55).


**Figure 22 Dual-redundant Meta SAN with 1606 Extension SAN Switch or MP Router**



**Figure 23 Dual-redundant Virtual Fabric**



**Figure 24 Dual-redundant VSAN**



25107b

## Supported routing configurations

Routing requires additional configuration rules for fabrics. For more information about routing configuration rules, see “Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules” (page 103), “MP Router fabric rules” (page 107), “C-series switches and fabric rules” (page 115), and “H-series switches and fabric rules” (page 131).


The typical configuration is a router or 8 Gb/s switch with integrated Fibre Channel routing connected to two or more fabrics, as shown in Figure 25 (page 52). An MP Router or 8 Gb/s switch with integrated Fibre Channel routing is required for fabrics that include B-series switches.

## Routing and core-edge fabrics

This section describes how to connect core switches when using routing to the following:

- MP Router for B-series Fibre Channel routing (Figure 25 (page 52))
- B-series Virtual Fabrics with IFR (Figure 26 (page 53))
- C-series VSANs with IVR (Figure 27 (page 53))

**Figure 25 1606 Extension SAN Switch or MP Router connecting core switches**



25108c

An alternative core switch routing configuration is to use the 8 Gb/s switch integrated Fibre Channel routing feature or MP Router Blades in the SAN Directors.

Figure 26 Virtual Fabrics connecting core switches





Figure 27 VSANs connecting core switches




## Routing through an IP network

When connecting fabrics through IP, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, MP Routers or Blades serve as FCIP gateways with Fibre Channel routing. Routers that communicate with the FCIP protocol must be installed in pairs (Figure 28 (page 54) and Figure 29 (page 54)).

HP supports FCIP configurations in which 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, 400 MP Router or MP Router Blade serve as an FCIP gateway and a Fibre


Channel switch. Servers and storage systems that support Continuous Access with FCIP can be directly connected to the Fibre Channel ports on the 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades (Figure 28 (page 54)), 400 MP Router or MP Router Blade (Figure 37 (page 59)).

**Figure 28 1606 Extension SAN Switches connecting fabrics through IP**



26598a

**Figure 29 MP Routers connecting fabrics through IP**

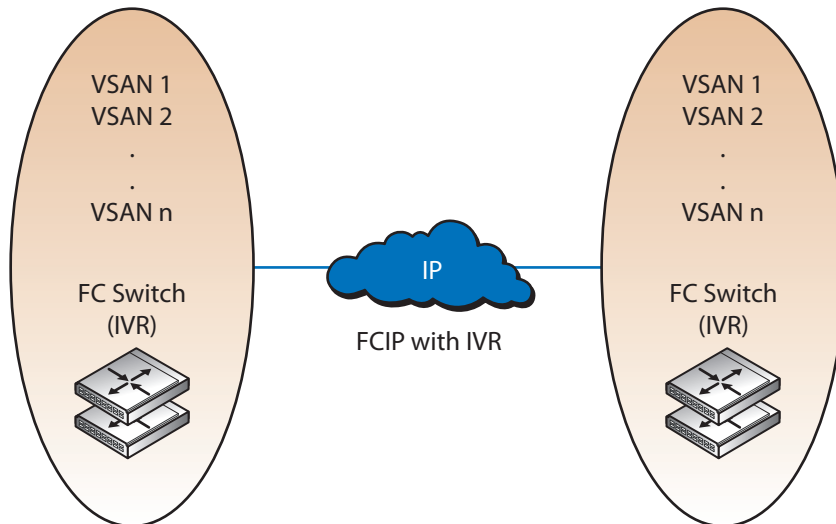


25152b

VSANs can be connected through IP using the FCIP functionality of the C-series Fibre Channel switches. See “C-series Fibre Channel switches” (page 115).

Figure 30 (page 54) shows IVR connecting VSANs through IP.

**Figure 30 IVR connecting VSANs through IP**



25111a

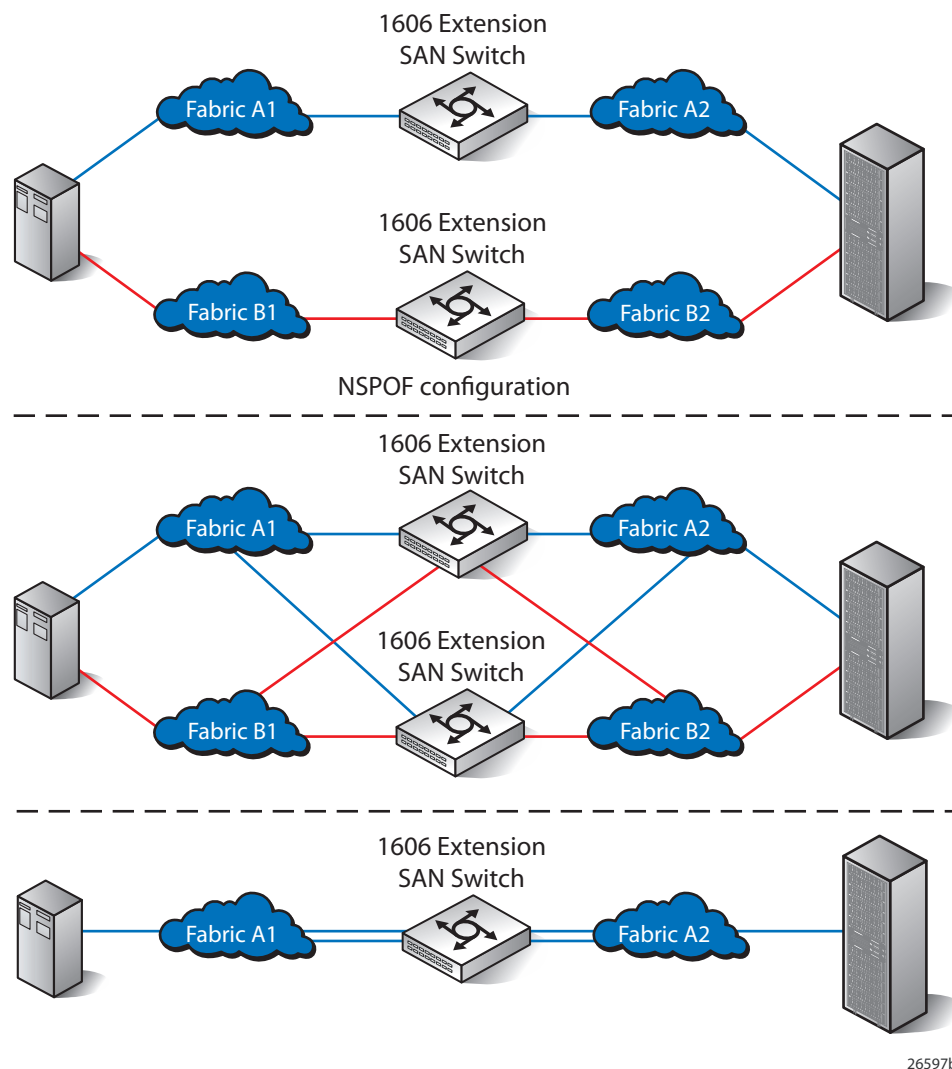
## High-availability router configurations

In high-availability configurations, use pairs of 8 Gb/s switches with integrated Fibre Channel routing, 1606 Extension SAN Switches, routers, or IFL pairs to provide redundant paths between fabrics. Figure 31 (page 55) and Figure 32 (page 56) show valid configurations.

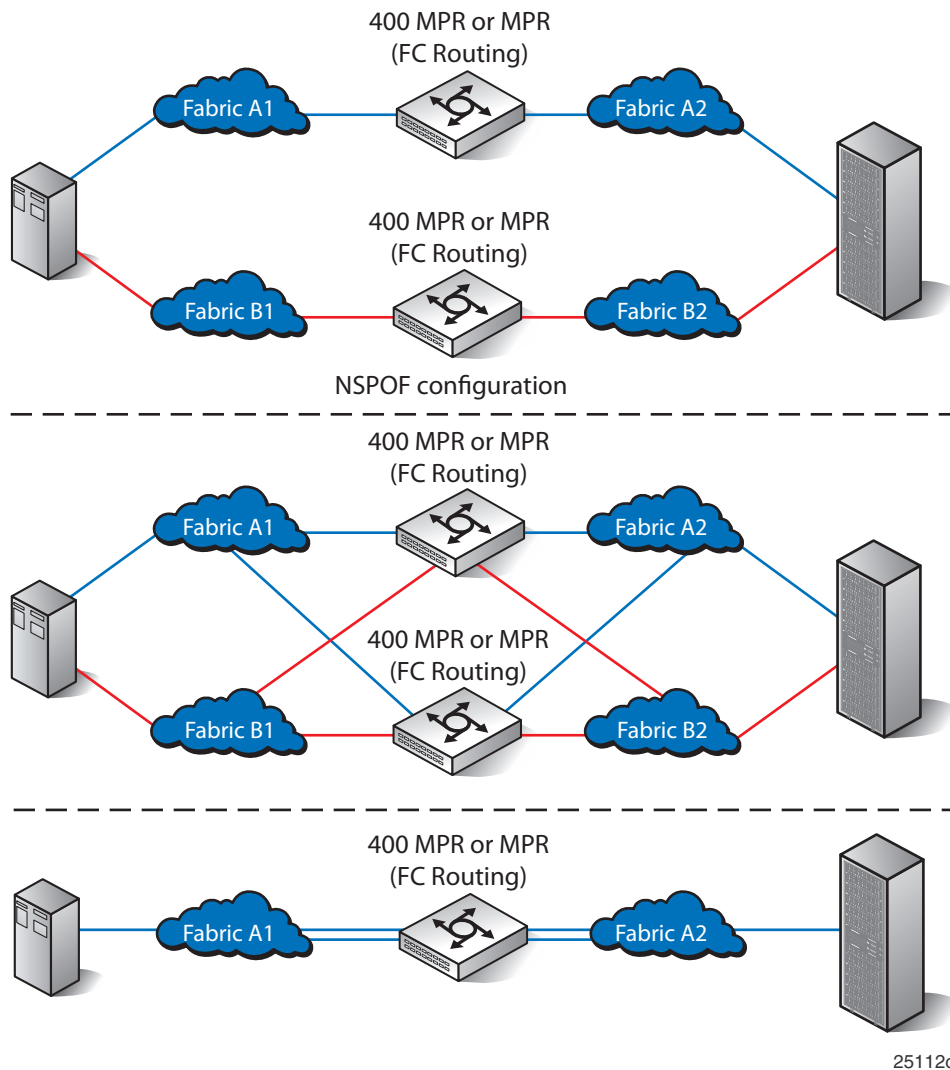
- The first configuration shows servers and storage connected using a pair of redundant fabrics in a level 4 NSPOF configuration. For information about high-availability levels, see “Data availability” (page 38).
- The second configuration shows routers cross-wired to provide full fabric connectivity in case a router fails.
- The third configuration shows multiple IFLs between the 1606 Extension SAN Switch or MP Router and edge fabrics to provide redundancy in case an IFL or MP Router port fails.

The configurations show a small number of SAN fabrics connected to each router. The same principles apply to configurations with a higher number of fabrics connected to a router, and a higher number of routers in a Meta SAN. For scalability rules, such as the maximum number of fabrics and MP Routers, see “B-series switches and fabric rules” (page 83).

**Figure 31 High-availability 1606 Extension SAN Switch configurations**



**Figure 32 High-availability MP Router configurations**



## 400 MP Router and MP Router Blade use cases

For configuration examples, see the 400 MP Router and MP Router Blade use-case white papers: <http://h18006.www1.hp.com/storage/whitepapers.html>

## H-series switches with TR configurations

The TR feature allows you to configure each 8 Gb port on the H-series switch as a TR\_Port, enabling the sharing of server and storage resources between fabrics through the industry-standard NPIV protocol. The TR feature also provides a higher level of security by ensuring that only customer-specified devices are visible and shared between fabrics.

Figure 17 (page 47) shows how an H-series switch can be in a fabric with other H-series switches and how you can connect it to multiple remote B-series or C-series fabrics. Figure 21 (page 50) shows how you can map each device connected to an H-series switch to multiple devices connected in one remote B-series or C-series fabric.



## Routing use cases

This section describes use cases for routing.

### SAN island consolidation and scaling

B-series routing consolidates SAN islands (multiple independent fabrics) into a Meta SAN. This modular SAN design offers:

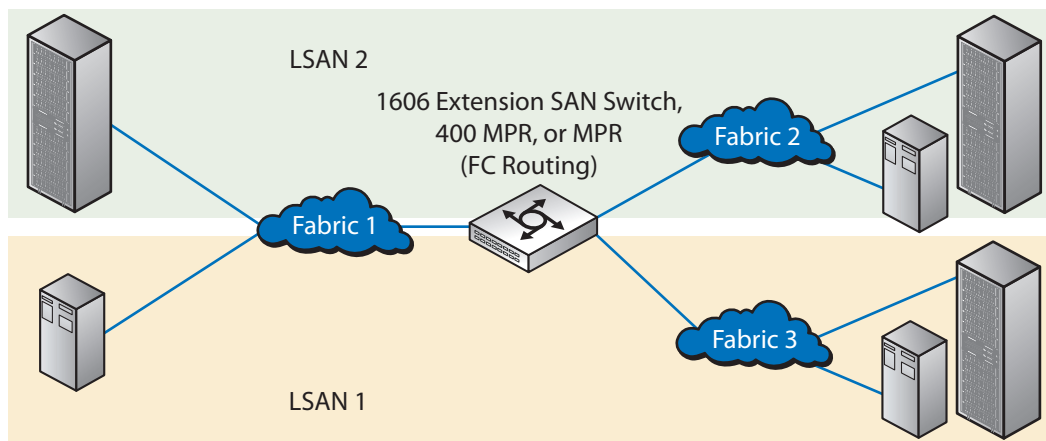
- Simplified scalability that allows you to scale a SAN without having to merge fabrics.
- Selective sharing of devices in different fabrics so that only devices required for specific functions are seen across fabrics.
- Limited sharing or specific times for data migrations and storage consolidation.
- Ability to access equipment without changing its physical location. Connecting multiple fabrics to 8 Gb/s switches with integrated Fibre Channel routing, 1606 Extension SAN Switches or DC Dir Switch MP Extension Blades, or MP Router or Blades enables sharing of devices located anywhere in the Meta SAN.
- Ability to connect B-series fabrics using Secure Fabric OS.
- Ability to connect B-series and M-series fabrics in a Meta SAN. For supported fabric modes and firmware versions, see “MP Router fabric rules” (page 107).

Routing does not merge fabrics, so existing zoning definitions and assigned domain IDs can be used without modification. Duplicate zoning definitions and domain IDs in fabrics are hidden by the MP Router. Fabrics in a Meta SAN can be scaled without affecting other fabrics.

Multiple SANs can be centralized and consolidated into one Meta SAN, or partitioned into different administrative domains or different logical fabrics using the Virtual Fabrics feature as required. HP recommends the use of Fabric Manager to simplify management procedures when implementing a Meta SAN.

Figure 33 (page 57) shows a typical configuration for SAN island consolidation.

**Figure 33 1606 Extension SAN Switch or MP Router consolidating SAN islands**



25113c

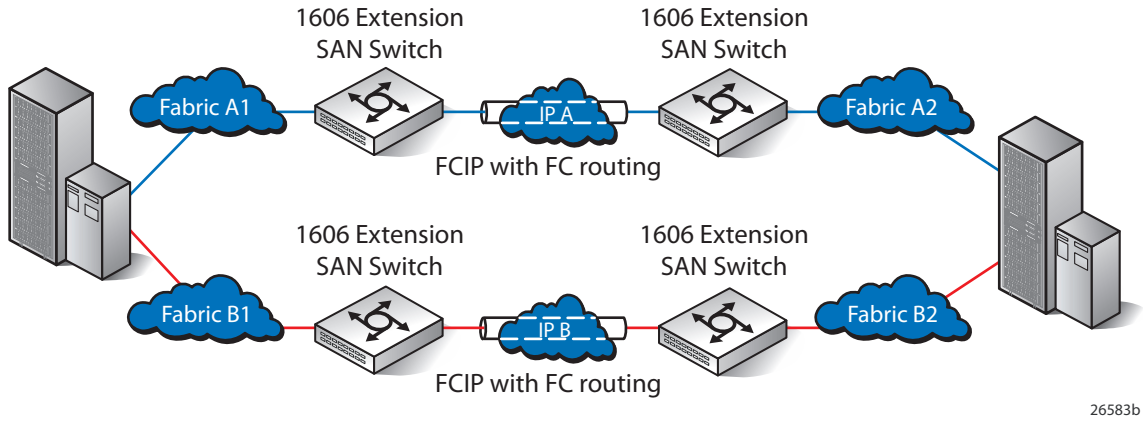
### Integration of Fibre Channel routing and FCIP

You can use the 1606 Extension SAN Switch or MP Router's integrated FCIP capability to extend disaster-tolerant applications such as HP Continuous Access for storage arrays.


In nonrouted FCIP configurations, local and remote fabrics merge when connected through an IP network. The IP connection acts as an ISL in a single fabric. By using Fibre Channel routing and FCIP, the local and remote fabrics connect without merging. You can create an LSAN that contains local and remote storage arrays and servers.

Figure 34 (page 58), Figure 35 (page 58), Figure 36 (page 58), Figure 37 (page 59), and Figure 38 (page 59) show typical HP Continuous Access NSPOF configurations in which the 1606 Extension SAN Switch or MP Router provides Fibre Channel routing and FCIP. Figure 36 (page 58) and Figure 37 (page 59) show a configuration in which the 1606 Extension SAN Switches or 400 MP Routers are used as both an FCIP gateway and Fibre Channel switch.


**Figure 34 NSPOF configuration with 1606 Extension SAN Switch providing Fibre Channel routing and FCIP with direct connect devices**




**Figure 35 NSPOF configuration with 400 MP Router providing Fibre Channel routing and FCIP**




**Figure 36 NSPOF configuration with 1606 Extension SAN Switch providing FC routing and FCIP with direct connect devices**



**Figure 37 NSPOF configuration with 400 MP Router providing Fibre Channel routing and FCIP with direct connect devices**



**Figure 38 NSPOF configuration with MP Router providing Fibre Channel routing and FCIP**




### Tape backup consolidation

The 1606 SAN Extension Switch or MP Router enables tape consolidation across multiple fabrics. Increased consolidation enables tape backup for devices in fabrics without tape libraries. Tape libraries and backup operations can be centralized and shared across multiple fabrics in a Meta SAN. There is no need to merge fabrics, which reduces equipment and management costs.

Figure 39 (page 60) shows a configuration in which a 1606 SAN Extension Switch or MP Router consolidates tape backup in a Meta SAN. For information about supported HP tape products and backup applications, see the *HP StorageWorks Enterprise Backup Solutions Design Guide* and the EBS compatibility matrix. These documents are available on the following websites:


- <http://www.hp.com/go/ebs>
- <http://h18004.www1.hp.com/products/storageworks/ebslegacymatrices.html>

**Figure 39 1606 Extension SAN Switch or MP Router consolidating tape backup in a Meta SAN**



Independent fabrics connected through a 1606 Extension SAN Switch or MP Router cannot have a direct ISL connection (Figure 40 (page 60)). A direct ISL connection between fabrics bypasses the MP Router, resulting in a full fabric merge.

**Figure 40 Unsupported 1606 Extension SAN Switch or MP Router configuration**



---

## 4 Fibre Channel over Ethernet

Fibre Channel over Ethernet is a protocol in which Fibre Channel frames can be sent over Ethernet networks. This allows Fibre Channel to use 10-GbE networks while preserving the Fibre Channel protocol.

The convergence of Ethernet and FC provides the same level of connectivity as each individual technology provides separately, but requires:

- 50% fewer switches in each server rack—Only two converged network ToR switches, compared with four switches per rack with separate Ethernet and FC switches (two network and two FC)
- 50% fewer adapters per server
- 75% fewer cable connections

This chapter describes the following topics:

- [“HP FCoE solutions overview” \(page 61\)](#)
- [“FCoE technology” \(page 62\)](#)
- [“HP FCoE products” \(page 63\)](#)
- [“FCoE configuration rules” \(page 69\)](#)
- [“Storage system support” \(page 78\)](#)

### HP FCoE solutions overview

HP offers multiple FCoE solutions:

- HP FCoE fabric-edge solutions deploy FCoE technology in Ethernet and FC environments with the benefit of convergence at the server and fabric edge using CNAs and CN switches. These solutions provide the benefits of both convergence and investment protection by allowing FCoE to be integrated with existing HP Fibre Channel SANs.
- HP end-to-end FCoE solutions for P6000/EVA storage provide an FCoE technology implementation from server to storage. These solutions consist of FCoE CNAs, CN switches, and either FCoE storage targets using the P6300/P6500 EVA with the FCoE target interface or the MPX200 Multifunction Router integrated with the P6000/EVA family of Fibre Channel storage systems. Three FCoE end-to-end configurations are available:
  - **P6300/P6500 EVA iSCSI/FCoE target**—Provides a 10-GbE FCoE and iSCSI interface on two ports of each controller. See [“P6300/P6500 EVA” \(page 316\)](#), and the iSCSI or iSCSI/FCoE configuration rules and guidelines chapter in the *HP P6300/P6500 Enterprise Virtual Array User Guide*.
  - **MPX200 Direct-connect**—Provides 10-GbE FCoE and iSCSI target connectivity for up to two P6000/EVA Fibre Channel storage systems without requiring any Fibre Channel switches.
  - **MPX200 Fabric-connect**—Provides additional FCoE and iSCSI storage connectivity, allowing for up to four P6000/EVA Fibre Channel storage systems using a fabric connection between the MPX200 and P6000/EVA storage.

The FCoE product set consists of the following:

- **HP A5820X-14XG-SFP+ Switch**—Provides a bridge function from a converged network to separate Ethernet and FC fabrics.
- **HP B-series and C-series CN switches**—Provide a bridge function from converged network to separate Ethernet and FC fabrics, or support end-to-end FCoE when used with CNAs and the P6300/P6500 EVA 10-GbE iSCSI/FCoE interface or MPX200 FCoE target functionality.

- **HP CN1000E, CN1100E, and CN1000Q CNAs**—Converge Ethernet and FC technologies in the server over 10-GbE links to converged network switches.
- **P6300/P6500 EVA storage with iSCSI/FCoE**—Provides native FCoE target connectivity within the storage system.
- **HP MPX200 Multifunction Router iSCSI/FCoE**—Provides FCoE target functionality when integrated with P6000/EVA Fibre Channel storage systems (for more information, see the *HP StorageWorks MPX200 Multifunction Router User Guide*).

## FCoE technology

The first generation of FCoE CN switches enable Ethernet and FC to coexist in a fabric and are designed as edge switches or ToR switches. Edge/ToR switches are typically deployed in redundant pairs installed at the top of a server rack. By using FCoE CN switches, you reduce the number of required switches by replacing separate Ethernet and FC switches with a converged edge/ToR switch. From the edge/ToR switch, ISL connections to the EoR FC and EoR Ethernet/IP switches provide access to the separate core layers of the FC fabric and Ethernet/IP network.

An FCoE blade that resides in a DC or DC04 SAN Director can also be used to integrate FCoE solutions with existing or new FC fabrics.

## Converged network switch ports

HP FCoE CN switches have two types of physical ports:

- 10-GbE ports
- 4-Gb or 8-Gb FC ports

FC ports can be used for either ISL connections to an existing FC fabric or connection directly to an HP storage array.

Table 10 (page 62) lists the number of ports for each HP FCoE CN switch.

**Table 10 Number of ports per switch**

HP switch	Number of 10-GbE ports	Number of FC ports
HP A5820 Converged Network Switch	14–18	4 to 8 (8 Gb/s)
2408 FCoE Converged Network Switch	24	8 (8 Gb/s)
DC SAN Director Switch 10/24 FCoE Blade	24	0 external ports 32 (8 Gb/s) backplane ports
C-series Nexus 5010 Converged Network Switch	20 to 26	0 to 8 (4 Gb/s) 0 to 6 (8 Gb/s)
C-series Nexus 5020 Converged Network Switch	40 to 52	0 to 16 (4 Gb/s) 0 to 12 (8 Gb/s)

**NOTE:** All FCoE CN switch models are by design CEE/IEEE DCB switches and are enabled for Ethernet by default. In order to be used for FCoE they must be enabled to forward FC frames to the FC ports, functioning as a Fibre channel forwarder.

The FCoE CN switch firmware update process causes momentary loss of Ethernet and Fibre Channel port connectivity. HP recommends that you implement a dual path configuration with redundant FCoE CN switches to minimize disruption during firmware updates.

## Interoperability

The HP A5820 Converged Network Switch interoperates with B-series, C-series, or H-series FC switches.

HP does not support interoperability between B-series and C-series CN and FC switches:

- 2408 FCoE Converged Network Switch and the 10/24 FCoE Blade operate with B-series FC switches only.
- C-series Nexus 5010/5020 Converged Network Switches operate with other C-series MDS FC switches only.

## HP FCoE products

This section describes the following HP FCoE CN switches and CNAs.

- “HP A5820 Ethernet/FCoE Switch” (page 63)
- “HP 2408 FCoE Converged Network Switch ” (page 64)
- “HP StorageWorks DC SAN Director Switch 10/24 FCoE Blade” (page 64)
- “HP C-series Nexus 5010/5020 Converged Network Switches” (page 65)
- “HP CN1000E and CN1100E” (page 66)
- “CN1000Q” (page 66)

## Converged network switches and blades

### HP A5820 Ethernet/FCoE Switch

The HP A5820 Switch Series features advanced flex-chassis switches that deliver a unique combination of unmatched 10 Gigabit Ethernet, Fibre Channel over Ethernet connectivity, high-availability architecture, full Layer 2/3 dual-stack IPv4/v6, and line-rate, low-latency performance on all ports.

Features and benefits :

- Powerful QoS features:
  - Creates traffic classes based on access control lists (ACLs), IEEE 802.1p precedence, IP, DSCP or Type of Service (ToS) precedence.
  - Supports filter, redirect, mirror, or remark.
  - Supports the following congestion actions: strict priority (SP) queuing, weighted round robin (WRR), weighted fair queuing (WFQ), weighted random early discard (WRED), weighted deficit round robin (WDRR), and SP+WDRR.
- Integrated network services—With support for open application architecture (OAA) modules, extends and integrates application capability into the network.
- Ring Resiliency Protection Protocol (RRPP)—Provides fast recovery for ring Ethernet-based topology; ensures consistent application performance for applications such as VoIP.
- 14 fixed 10-GbE ports
- Two expansion bays, one 4-port FCoE expansion module required. Second bay can be used for second FCoE module or additional 10-GbE ports.
- Optional expansion modules:
  - HP 4-port 8/4/2 Gb/s FCoE SFP+ A5820 Module
  - HP A5820X/A5800 4-port 10-GbE SFP+ Module
  - HP A5820X/A5800 2-port 10-GbE SFP+ Module

- Compatible with B-series, C-series, and H-series FC switches
- Considerations:
  - Requires connectivity to an FC switch for access to FC storage systems. See [Figure 42 \(page 69\)](#) and [Figure 44 \(page 71\)](#).
  - Access to native FCoE target devices not supported via 10-GbE ports

## HP 2408 FCoE Converged Network Switch

The 2408 FCoE Converged Network Switch is a rebranded version of the Brocade 8000 FCoE Switch.

### Features

- Next-generation Ethernet L2/L3 switch and next-generation FC switch merged into one product
- CEE support
  - Full industry-standard implementation
  - Supports FCoE and FCoE Initialization Protocol (FIP)
- 24 10-GbE and 8 8-Gb/s FC ports
- Uses 1 RU of space
- Leverages existing Brocade fabric operating system (FOS)
- Supports link aggregation (LAG) on DCB ports
- Supports Brocade ISL trunking on FC ports

### Considerations

- The port types are fixed. You cannot use 10-GbE ports for FC connections, and you cannot use FC ports for 10-GbE connections.
- L3 routing features are not currently supported.
- 10-GbE ports support virtual F\_Ports only (virtual E\_Ports are not supported).

## HP StorageWorks DC SAN Director Switch 10/24 FCoE Blade

The HP DC SAN Director Switch 10/24 FCoE Blade is a rebranded version of the Brocade FCOE10-24 blade.

### Features

- 24 10-GbE ports
- 32 8-Gb/s FC ports on the backplane (DC and DC04 SAN Directors only)
- Hot pluggable
- Blade power and status LEDs
- Link status LEDs for each port
- FCoE switching
- CEE switching
- L2 Ethernet protocols STP/MSTP/RSTP, VLAN tagging, and link aggregation
- Standard Ethernet encapsulation



## Considerations

- The DC SAN Director Switch 10/24 FCoE Blade has 24 FCoE ports and can be installed in either a DC or DC04 SAN director.
- The DC SAN Director Switch 10/24 FCoE Blade enables server edge connectivity by connecting CNAs directly to any of its 24 Ethernet ports.
- Storage can be connected to:
  - Any other Fibre Channel blade in the same director
  - Any Fibre Channel switch that is in the same fabric as the DC or DC04 SAN Director that contains the DC SAN Director Switch 10/24 FCoE Blade
- The DC SAN Director Switch 10/24 FCoE Blade supports optical cabling and SFP+ transceivers only.
- There are no licensing requirements for this blade.

## HP C-series Nexus 5010/5020 Converged Network Switches

HP resells the Cisco Nexus 5010 and 5020 Switches.

### Features

- L2 access:
  - IEEE DCB and FCoE support  
IEEE DCB is a Cisco unified fabric product with additional proprietary features.
  - NX-OS with combined features from Cisco IOS and Cisco MDS 9000 SAN-OS/NX-OS
- Cisco 5020 (2 RU):
  - 40 10-GbE ports
  - Two optional expansion module slots
  - Up to 52 10-GbE ports, or a combination of 40 10-GbE ports and 16 FC ports
- Cisco 5010 (1 RU):
  - 20 10-GbE ports
  - One optional expansion module slot
  - Up to 26 10-GbE ports, or a combination of 20 10-GbE ports and 8 FC ports
- Optional expansion modules:
  - 6-port 10-GbE expansion module
  - 6-port 8-Gb/s FC expansion module
  - 8-port 4-Gb/s FC expansion module
  - 4-port 4-Gb/s FC and 4-port 10-GbE expansion module

## Converged network adapters

Servers with CNAs require fewer adapters per server and, therefore, 75% fewer cable connections for network and FC attachments.

HP offers the HP CN1000E CNA, the HP CN1100E CNA, and the HP CN1000Q CNA. For information about servers that support CNAs, see [“Server support” \(page 74\)](#).

[Table 11 \(page 66\)](#) lists the supported HP CNAs.

**Table 11 HP CNA part numbers**

HP part number	HP model number	Description
BK835A	CN1100E	Dual port converged network adapter (no cable or transceiver is included)
AW520A	CN1000E	
BS668A	CN1000Q	

## HP CN1000E and CN1100E

### Features

The HP CN1000E and CN1100E CNAs have the following features:

- Ships with half-height and full-height brackets
- Dual ports for redundancy
- Full 10-Gb/s bandwidth on both ports
- Each port can operate as a NIC and/or FCoE port
- 2 SFP+ connectors
- Supports optical or copper cables

### Considerations

- x8 PCI Express Gen2 card
- Requires 14.5 W of power
- 1 GbE is not supported

## CN1000Q

### Features

The HP CN1000Q and CN1100E CNAs have the following features:

- Ships with half-height and full-height brackets
- Dual ports for redundancy
- Full 10-Gb/s bandwidth on both ports
- Each port can operate as a NIC and/or FCoE port
- 2 SFP+ connectors
- Supports optical or copper cables

### Considerations

- x8 PCI Express Gen2 card
- Requires 19 W of power
- 1 GbE is not supported

## FCoE storage systems

HP supports the following FCoE storage system:

- P6300/P6500 EVA with iSCSI/FCoE 10-GbE front end host ports (2 per controller)

### iSCSI/FCoE module

The iSCSI/FCoE module is configured in a dual-controller configuration in the HP P6300/P6500 EVA (see [Figure 41 \(page 68\)](#)). Dual-controller configurations provide for high availability with failover between iSCSI/FCoE modules. All configurations are supported as redundant pairs only. iSCSI connected servers can be configured for access to one or both controllers. For information about P6300/P6500 EVA Fibre Channel support, see “[P6000/EVA storage system rules](#)” (page 213). For information about P6300/P6500 EVA 1-GbE iSCSI support, see “[iSCSI storage](#)” (page 312).

### Operating system and multipath software support

For the latest information on operating system and multipath software support, see the HP SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

[Table 12 \(page 67\)](#) provides the operating system and multipath software support for P6300/P6500 EVA 10-GbE iSCSI/FCoE.


**Table 12 P6300/P6500 EVA 10-GbE iSCSI/FCoE operating system and multipath software support**

Operating system	Multipath software		Clusters EVA storage system
Microsoft Windows Server 2008	MPIO with HP DSM (supports both CN1000E and CN1100E)	MPIO with Microsoft DSM (supports CN1000E only)	Failover Clustering, MSCS
Microsoft Windows Server 2003	MPIO with HP DSM (supports CN1000E only)	MPIO with Microsoft DSM (supports CN1000E only)	
Red Hat Linux, SUSE Linux	Device Mapper		None
VMware	VMware MPxIO		None

For more information, see the iSCSI or iSCSI/FCoE configuration rules and guidelines chapter in the *HP P6300/P6500 Enterprise Virtual Array User Guide*.

[Figure 41 \(page 68\)](#) describes the P6300/P6500 EVA 10-GbE iSCSI/FCoE configuration.

Figure 41 P6300/P6500 EVA iSCSI/FCoE end-to-end configuration




26659a

# FCoE configuration rules

This section describes configuration rules for the HP FCoE solution. The rules are defined for the CNAs and each of the FCoE CN switch series.

Figure 42 (page 69) shows a converged SAN fabric with servers using CNAs only.

**Figure 42 Converged SAN fabric configuration**



26647c

Figure 43 (page 70) shows a converged SAN fabric with servers using CNAs connected to a B-series CN switch and a DC SAN Director Switch 10/24 FCoE Blade.

**Figure 43 Converged SAN fabric with an FCoE blade**

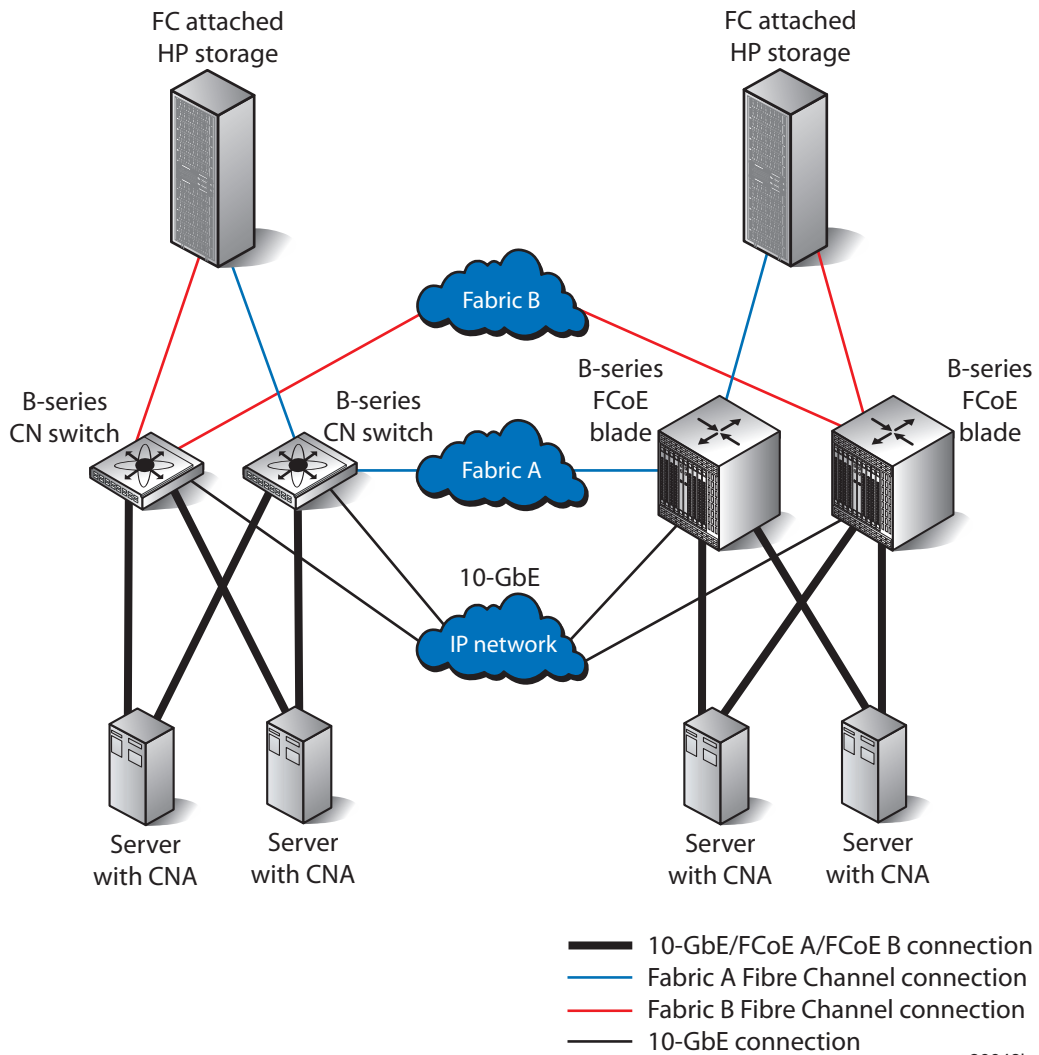


Figure 44 (page 71) shows a converged fabric that has servers with CNAs and servers with FC HBAs and NICs.

Figure 44 FCoE integrated with FC SAN fabric





Figure 45 (page 72) shows an FCoE end-to-end direct-connect storage configuration using an MPX200 Multifunction Router and P6000/EVA Fibre Channel storage systems. This provides 10-GbE FCoE and iSCSI connectivity for up to two P6000/EVA Fibre Channel storage systems without requiring any Fibre Channel switches.

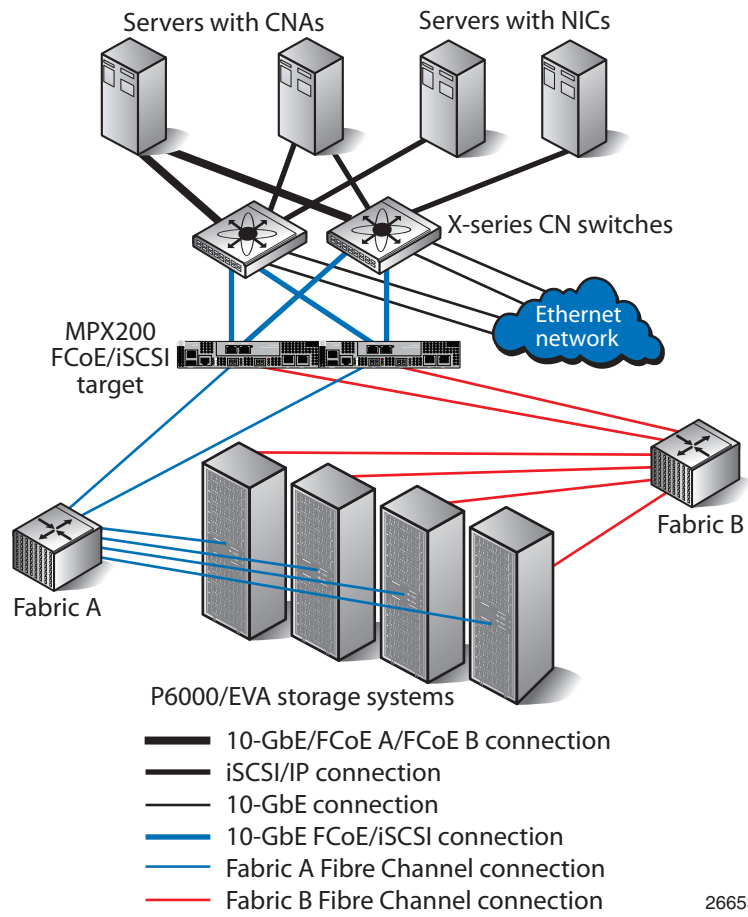
**Figure 45 FCoE end-to-end direct-connect P6000/EVA Fibre Channel storage configuration**



The configuration shown in Figure 46 (page 73) provides additional connectivity, allowing for up to four P6000/EVA storage systems using a fabric connection between the MPX200 and P6000/EVA storage. In addition to FCoE and iSCSI, Fibre Channel connected servers can access the same P6000/EVA storage systems through the Fibre Channel fabric. This configuration provides the highest level of multi-protocol support with FCoE, iSCSI, and Fibre Channel connectivity for up to four P6000/EVA storage systems.



**Figure 46 FCoE end-to-end fabric-connect P6000/EVA storage configuration**



## CNA configuration rules

**Table 13 HP CNA requirements**

Item	Description
Number of CNAs per server (maximum)	2
Cable types and distances	<ul style="list-style-type: none"><li>• SFP+ optical CN1000E: 40m, CN1100E and CN1000Q: 100m</li><li>• Copper 7m and 10m (<b>Note:</b> C-series FCoE CN switches only)</li></ul>

For current firmware and driver support, see the HP SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

### Server support

FCoE is supported with ProLiant G5, G6, and G7 servers.

For current server support, see the server QuickSpecs at <http://www.hp.com/go/quickspecs/>.

### HP A5820 Converged Network Switch

This section describes support for the A5820X-14XG FCoE CN switch and FCoE fabric rules. For more information on configuring the A5820X-14XG FCoE CN switch, see “FCoE switch configuration quick-setup instructions” (page 379).

### Usage

The A5820X switch with its FCoE module is intended for use as a server edge switch, providing a bridge between 10-GbE and Fibre Channel protocols. See “FCoE fabric topologies” (page 38).

### FCoE Converged Network Switch fabric rules

This section describes the fabric rules for the A5820X FCoE CN switch and other factors you should consider when building A5820X FCoE CN fabrics. For Fibre Channel switch fabric rules pertaining to the specific family of FC switches being used, see “Fabric infrastructure rules” (page 82).

**Table 14 HP A5820X CN switch fabric rules**

Rule number	Description
1	At least one HP 4-port 8/4/2 Gb/s FCoE SFP+ A5820 module is required for the A5820X-14XG to operate as an FCoE CN switch.
2	At least one FC port of the A5280 FCoE module must be connected to a B-series, C-series, or H-series FC switch.
3	NPIV must be enabled on all FC switches connected to the A5820 FCoE module.
4	The A5280 FCoE module FC ports only support FC switch connections.
5	Native FCoE storage is not supported using the A5820 10-GbE ports.

**Table 14 HP A5820X CN switch fabric rules** (continued)

Rule number	Description
6	<p>Up to two A5820 FCoE modules are supported in a single A5820X switch. However, when using two A5280 FCoE modules, consider the following:</p> <ul style="list-style-type: none"> <li>When two A5820 FCoE modules are connected to the same SAN, the FCoE modules attempt to load balance the server load across the two FCoE modules as servers are added by distributing the CNA logins between the two modules. Failover from one module to the other is supported in this configuration, but failback is not supported.</li> <li>The current Emulex CNA firmware does not support the load balancing method used when two A5820 FCoE modules are connected to the same SAN, so the distribution of the CNA logins may not be optimal. This issue will be corrected in a future Emulex CNA firmware release.</li> </ul>
7	<p>Connection to Brocade 8 Gb/s switches and ARB(FF)/ARB(FF) mode setting: HP recommends all 8 Gb/s Brocade switches running FOS V6.3.1a and later be set to mode 3 using the <code>portcfgfillword</code> command. For more information, see the <i>Fabric OS Command Reference Manual</i> supporting FOS V6.4 and the FOS release notes.</p>
8	<p>When connecting the A5820 to multiple SANs, unique discovery and data VLANs are required for each SAN.</p>

**Firmware, CNA, operating systems, and storage products support**

For current information on supported A5820X firmware, operating systems, and options, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

**B-series FCoE Converged Network switch**

This section describes support for B-series FCoE CN switches, switch model, and FCoE fabric rules. For more information on configuring B-series FCoE CN switches, see “FCoE switch configuration quick-setup instructions” (page 379).

**NOTE:** The naming scheme for the HP StorageWorks 2408 Converged Network Switch, which has 24 CEE ports and 8 Fibre Channel ports, is slightly different from the typical B-series switch naming scheme.

Table 15 (page 75) describes the B-series FCoE CN switch.

**Table 15 B-series FCoE CN switch**

Switch name	Firmware version		Fabric Management		Number of ports
			FM <sup>1</sup>	DCFM <sup>2</sup>	
HP StorageWorks 2408 FCoE Converged Network Switch	7.0.0a		N/A		24 CEE ports 8 FC ports
	6.4.2a	6.4.0c			
	6.4.1b	6.3.2b		10.4.5	
	6.4.1a	6.3.2b		10.3.4 <sup>3</sup>	
	6.4.1	6.3.2a			

<sup>1</sup> FM = Fabric Manager

<sup>2</sup> DCFM = Data Center Fabric Manager

<sup>3</sup> DCFM only provides the discovery function for the FCoE switch, and launches WebTools for switch management with DCFM versions prior to 10.3.0. Full management features are provided with DCFM 10.3.0 and later versions.

For the latest information on supported B-series FCoE CN switches and firmware versions, see the HP SAN Infrastructure website: <http://h18006.www1.hp.com/storage/saninfrastructure.html>.

## Features

Features of the 2408 FCoE Converged Network Switch are the same as the B-series 8Gb switches, but the following features are not supported:

- Advanced Zoning (port-based and QoS)
- Extended Fabrics
- Adaptive Networking with QoS
- Virtual Fabrics
- Integrated Fibre Channel routing
- FICON
- Admin Domains
- TI Zones
- M-EOS Interop
- Nondisruptive firmware upgrade

Table 16 (page 76) provides a comparison of the high-availability features for the B-series FCoE CN switch.

**Table 16 B-series FCoE CN switch high-availability feature comparison**

Switch model	Redundant /hot-swappable power	Redundant /hot-swappable cooling	Redundant core switching blade	Redundant control processor	Non-disruptive code activation	Non-disruptive port expansion	Redundant active components
HP StorageWorks 2408 FCoE Converged Network Switch	Yes/Yes	Yes/Yes	No	No	No	N/A	No

The HP StorageWorks DC SAN Director Switch 10/24 FCoE Blade (DC SAN Director Switch 10/24 FCoE Blade) provides the following features:

- 24 FCoE ports operating at 10 Gb/s
- 32 FC ports operating at 8 Gb/s through the backplane
- Hot pluggable
- Blade power and status LEDs
- Port link-status LEDs
- FCoE switching
- CEE switching
- L2 Ethernet protocols STP/MSTP/RSTP, 802.1q; and Link Aggregation (802.1ad)
- Standard Ethernet encapsulation

## Operating systems and storage products

The fabric rules for the 2408 FCoE Converged Network Switch apply to SANs that include the operating systems and storage products listed in Table 20 (page 79).

## C-series FCoE Converged Network switches

This section describes support for C-series FCoE CN switches, switch models, and FCoE fabric rules. For more information on configuring C-series FCoE CN switches, see “FCoE switch configuration quick-setup instructions” (page 379).

## C-series Nexus 5020 FCoE Converged Network Switch

The C-series Nexus 5020 is a 2 RU, 10-GbE, Cisco IEEE DCB, FCoE, and Fibre Channel switch. It has a fixed configuration of 40 10-GbE, Cisco IEEE DCB, and FCoE SFP+ ports. The first 16 fixed ports are dual speed, supporting both 10-GbE and 1-GbE. The switch has two expansion slots that can accommodate up to 16 Fibre Channel switch ports.

## C-series Nexus 5010 FCoE Converged Network Switch

The C-series Nexus 5010 is a 1 RU, 10-GbE, Cisco IEEE DCB, FCoE, and Fibre Channel switch. It has a fixed configuration of 20 10-GbE, Cisco IEEE DCB, and FCoE SFP+ ports. The first 8 fixed ports are dual speed, supporting both 10-GbE and 1-GbE. The switch has one expansion slot that can accommodate up to 8 Fibre Channel switch ports.

## Expansion modules for the C-series Nexus 5000 Series switches

The C-series Nexus 5000 Series switches can accommodate expansion modules that connect Fibre Channel SANs to 8, 4, 2, or 1 Gb/s Fibre Channel switch ports. Use these modules to increase the number of 10-GbE, Cisco IEEE DCB, and FCoE ports. The Nexus 5010 has one expansion slot; the Nexus 5020 has two expansion slots that accommodate the following modules:

- Nexus 5000 6-port 10GbE Module
- Nexus 5000 4-port 4Gb FC and 4-port 10GbE Module
- Nexus 5000 8-port 4Gb FC Module
- Nexus 5000 6-port 8Gb FC Module

For the latest information on supported C-series FCoE CN switches and firmware versions, see the HP SAN Infrastructure website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

Table 17 (page 77) lists the firmware version and port maximums for the C-series Nexus switches.

**Table 17 C-series Nexus switches for NX-OS**

Switch	Firmware version	Maximum number of Fibre Channel ports	Maximum number of Ethernet ports
Nexus 5020	NX-OS 5.0(2)N2(1), 4.2(1)N2(1), 4.2(1)N1(1), 4.1(3)N2(1a), 4.1(3)N1(1), 4.0(1a)N1(1a), 4.0(1a)N1(1)	16	52
Nexus 5010		8	26

Table 18 (page 77) describes C-series Nexus expansion module support.

**Table 18 C-series Nexus expansion module support matrix**

Expansion module	Switch	
	Nexus 5020	Nexus 5010
HP Nexus 5000 6-port 10GbE Module	Yes	Yes
HP Nexus 5000 4-port 4Gb FC and 4-port 10GbE Module	Yes	Yes
HP Nexus 5000 8-port 4Gb FC Module	Yes	Yes
HP Nexus 5000 6-port 8Gb FC Module	Yes	Yes

## Features

Features of the C-series FCoE CN switches include:

- Unified fabrics
- I/O consolidation—Reduces capital and operating expense
- Application flexibility—Support a number of applications as access-layer switches co-located with servers in data center racks or for middle-of-row deployments
- Consolidated resource management—Support existing configuration and management tools
- Energy efficiency
- PFC
- Bandwidth management
- DCBX Protocol
- Congestion management (BCN)
- Layer 2 multipathing (hardware capable)
- Convergence and compatibility with existing FC fabrics
- High availability (see [Table 19 \(page 78\)](#))

[Table 19 \(page 78\)](#) provides a comparison of the high-availability features for C-series Nexus 5000 Converged Network Switches.

**Table 19 C-series Nexus 5000 Series switch high-availability feature comparison**

Switch	Redundant/ hot-swappable power	Redundant/ hot-swappable cooling	Redundant control processor	Nondisruptive code activation	Expansion module support	Protocol support
Nexus 5020	Yes/Yes	Yes/Yes	No	No	Yes	FC
Nexus 5010	Yes/Yes	Yes/Yes	No	No	Yes	IEEE DCB 10-GbE FCoE

## Usage

The Nexus 5000 Series switches are intended for use as edge switches, providing a bridge between 10-GbE and Fibre Channel protocols. See [“FCoE fabric topologies” \(page 38\)](#).

## FCoE Converged Network switch fabric rules

This section describes the fabric rules for C-series FCoE CN switches and other factors you should consider when building C-series FCoE CN fabrics. For general C-series Fibre Channel switch fabric rules, see [“Fibre Channel switch fabric rules” \(page 126\)](#).

## Operating systems and storage products

The fabric rules for C-series FCoE CN switches apply to SANs that include the operating systems and storage products listed in [Table 20 \(page 79\)](#).

## Storage system support

This section describes storage system support for access from CNA-based servers. These storage systems can be attached to Fibre Channel switches in the fabric or connected to the Fibre Channel ports of an FCoE CN switch.

SAN boot is not supported when storage systems are connected to the Fibre Channel ports of an FCoE CN switch. SAN boot is supported for storage attached to the Fibre Channel ports in a SAN Director that contains a 10/24 FCoE blade and based on the current support listed for Fibre Channel switches.

Table 20 (page 79) lists operating system software and storage products that are supported with HP CN1000E and CN1000Q CNAs.

**Table 20 HP CN1000E and CN1000Q operating system and storage system support**

Operating systems <sup>1</sup>	Storage products <sup>2</sup>
<ul style="list-style-type: none"> <li>• Microsoft Windows</li> <li>• Linux</li> <li>• VMware ESX</li> </ul>	<ul style="list-style-type: none"> <li>• MSA2000fc G2 (MSA2312fc/MSA2324fc)</li> <li>• MSA2000fc (MSA2012fc/MSA2212fc)—Supported on CN1000E only</li> <li>• P2000 G3</li> <li>• EVA4400</li> <li>• EVA4400 with embedded switch<sup>3</sup></li> <li>• EVA4100/6100/8100</li> <li>• EVA6400/8400</li> <li>• XP10000/12000</li> <li>• XP20000/24000</li> <li>• 3PAR<sup>4</sup> F-Class, T-Class</li> <li>• P6300/P6500 EVA</li> <li>• P6300/P6500 EVA with an iSCSI/FCoE module<sup>5</sup></li> </ul>

<sup>1</sup> For operating system and storage product version support, see the HP SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

<sup>2</sup> Supported as Fibre Channel targets unless noted otherwise.

<sup>3</sup> The EVA4400 with the embedded switch is not supported for connectivity to C-series Fibre Channel switches or Nexus Converged Network Switches.

<sup>4</sup> For information about 3PAR CNA and FCoE CN switch support, see the *3PAR InForm Compatibility Matrix*.

<sup>5</sup> Supported as an FCoE/iSCSI target connected to a CN switch 10 GbE Ethernet port.

## Operating system support

For current operating system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## Boot from SAN support

BFS is not supported for FC arrays attached to 2408 FCoE Converged Network Switches or C-series Nexus 5000 Converged Network Switches at this time. BFS is supported for FC arrays attached to standard FC switches in a fabric that includes CN switches.

DC SAN Director Switch 10/24 FCoE Blade supports BFS from the storage attached to any of the other FC blades in the same director or to any Fibre Channel switch that is part of the same fabric as the DC SAN Director that contains the 10/24 FCoE Blade.

This section provides information about installing BFS software.

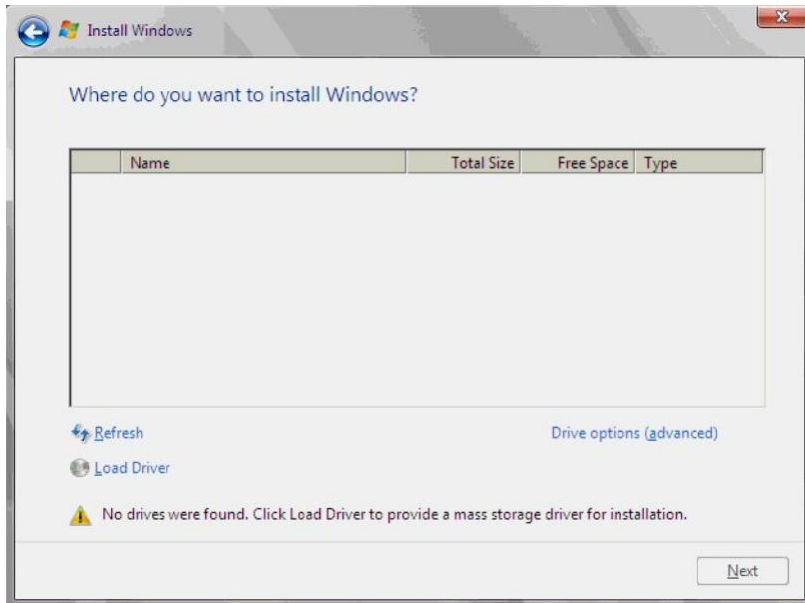
### Windows 2003 (x86, x64) and 2008 (x86, x64)

- BFS is supported on all supported HP servers.
- The information in *HP Boot from SAN Configuration Guide* (available on <http://h18006.www1.hp.com/storage/networking/bootsan.html>) is accurate for Windows BFS on a CNA based server for Windows 2003 with the following exceptions:
  - You must install an HP CNA driver kit (available for download at the HP website:

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

- For Windows 2003, press **F6** to prompt for the driver disk.
- For Windows 2008, click **Load Driver** when prompted for the location in which to install Windows (see Figure 47 (page 80)).

**Figure 47 Install Windows 2008 (x86, x64)**



## Red Hat Enterprise Linux

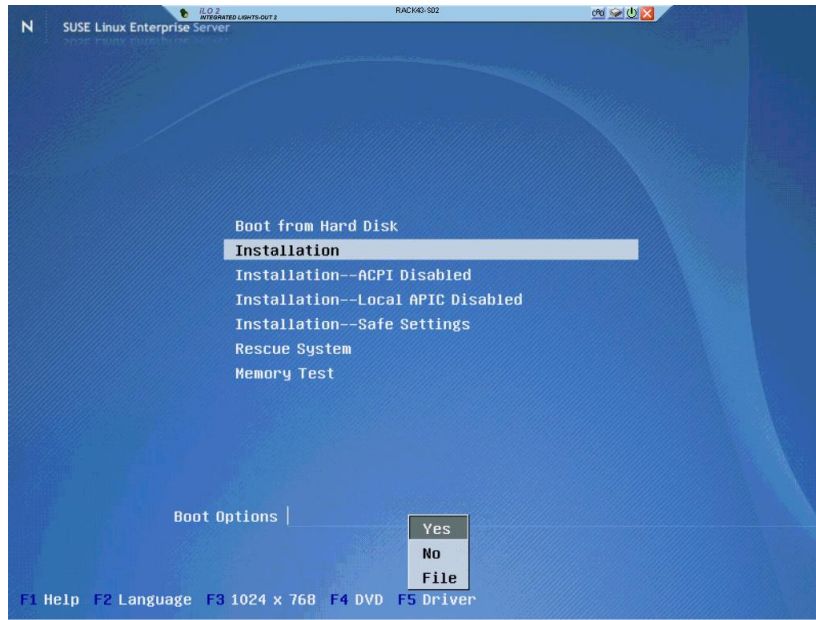
- BFS is supported on all supported HP servers.
- The information in *HP Boot from SAN Configuration Guide* (available at <http://h18006.www1.hp.com/storage/networking/bootsan.html>) is accurate for BFS on a CNA based server with the following exceptions:
  - You must install an HP CNA driver kit (available for download at the HP website: <http://www.hp.com/support/downloads>).
  - Use the `linux dd mpath` rather than the `linux mpath` command to prompt to load a driver and enable multipath for the installation.



## SUSE Linux Enterprise Server

- BFS is supported on all supported HP servers.
- The information in *HP Boot from SAN Configuration Guide* (available at <http://h18006.www1.hp.com/storage/networking/bootsan.html>) is accurate for BFS on a CNA based server with the following exception:
  - Use the GUI to load the driver (Figure 48 (page 81)).

**Figure 48 SLES installation**



## Multipathing software support

For current multipathing software support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

---

## Part II Fabric infrastructure rules

Fabric infrastructure rules are presented in these chapters:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “H-series switches and fabric rules” (page 131)
- “M-series switches and fabric rules” (page 140)
- “SAN fabric connectivity and switch interoperability rules” (page 146)

---

## 5 B-series switches and fabric rules

This chapter describes the B-series Fibre Channel and FCoE switches and blades; Extension SAN Switches and Blades; MP Routers, and fabric rules for building B-series fabrics. It describes the following topics:

- “B-series Fibre Channel switches, 1606 Extension SAN Switches, and MP Routers” (page 83)
- “Fibre Channel switch fabric rules” (page 94)
- “B-series FCoE Converged Network switch” (page 75)

---

**NOTE:** For information about using switches from the different series in the same SAN or the same fabric, see “SAN fabric connectivity and switch interoperability rules” (page 146).

---

### B-series Fibre Channel switches, 1606 Extension SAN Switches, and MP Routers

B-series Fibre Channel switches:

- Can be core or edge switches. When configured in a core-edge fabric topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- Have 8, 16, 24, 32, 40, 48, 64, 80, 128, 192, 256, 384, 512 ports.
- Include entry-level and high-availability switches.
- Support plug-and-play compatibility.
- Support 10 Gb/s ISL connectivity (through a 10/6 FC ISL Blade for DC SAN Backbone Director, DC04 SAN Director, or 4/256 SAN Director, or by configuration of a 16 Gb port for 10 Gb/s on 16 Gb switch or director blade ISL ports).
- Support Fibre Channel routing. See “1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules” (page 100), “Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules” (page 103), and “MP Router fabric rules” (page 107).
- Support FCIP through the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, MP Router Blade, MP Router, and HP IP Distance Gateway. See “SAN extension” (page 249).
- Support iSCSI through the MP Router and iSCSI Director Blade. See “iSCSI storage” (page 312).
- Support data at rest encryption through the encryption switch and encryption blade.
- Support in-flight encryption and/or compression on 16 Gb switch or director blade ISL ports.

Switch blade support is as follows:

- HP SN8000B 8-Slot SAN Backbone Director and HP SN8000B 4-Slot SAN Director support the 64-port 8 Gb/s Fibre Channel blade, 32 and 48-port 16 Gb/s Fibre Channel blades, encryption blade, and DC Dir Switch MP Extension Blade.
- DC SAN Backbone Director and DC04 SAN Director support the 16, 32, 48, and 64-port 8 Gb/s Fibre Channel blades, 10 Gb/s 10/6 FC ISL Blade, MP Router Blade, encryption blade, and the DC Dir Switch MP Extension Blade.
- SAN Director 4/256 supports the 16, 32, and 48-port 4 Gb/s Fibre Channel blades, 16, 32, and 48-port 8 Gb/s Fibre Channel blades, 10 Gb/s 10/6 FC ISL Blade, the MP Router Blade, and iSCSI Director Blade.

The B-series switches offer:

- High availability
- Scalability
- Cost efficiency

## Model numbering

The B-series Fibre Channel switches typically use the numbering scheme  $x/y$ :

- $x$ —The highest speed at which the switch ports can operate, measured in Gb/s
- $y$ —The number of switch ports

For example, the HP StorageWorks SAN Switch 4/16 is a 4 Gb/s switch with 16 ports.

The 16 Gb/s, 8 Gb/s, 4 Gb/s, and 2 Gb/s switch ports autonegotiate the signaling speed. For example, when you connect a 4 Gb/s or 2 Gb/s port directly to a 1 Gb/s port, both ports run at 1 Gb/s in each direction. If the ports are not directly connected, the fabric switch that connects the ports determines the signaling speed.

## Model naming

The B-series Fibre Channel switches are typically called HP SAN, SAN Backbone, Core, or Director switches. SAN switches are edge (entry-level or mid-range) switches. SAN Backbone, Core, and Director switches are core (enterprise-class) switches. The suffix -EL, V, or N identifies entry-level switches.

The B-series embedded switches for HP c-Class BladeSystem servers are called the Brocade 8Gb SAN Switch for HP BladeSystem c-Class and the Brocade 4Gb SAN Switch for HP c-Class BladeSystem. The embedded switch for HP p-Class BladeSystem servers is called the Brocade 4Gb SAN Switch for HP p-Class BladeSystem. The B-series Fibre Channel routers are called the HP StorageWorks 400 Multi-protocol Router (400 MP Router), Multi-protocol Router Blade (MP Router Blade), and Multi-protocol Router (MP Router). The 8 Gb SAN extension products are called the HP StorageWorks 1606 Extension SAN Switch and HP StorageWorks DC SAN Director Multi-protocol Extension Blade.

The new HP model naming convention has the B-series Storage Networking switches named SNxxxxB where SN stands for Storage Networking switch and B for B-series. The xxxx numbering is 8000 (Director Class) or 6000 (Fabric Class).

## Switch models

Table 21 (page 85) and Table 22 (page 86) describe the B-series Fibre Channel switches and routers currently available from HP (16 Gb/s and 8 Gb/s B-series Fibre Channel switches).

Table 23 (page 87) describes the legacy Fibre Channel switches and routers supported by HP.

HP supports all B-series Fibre Channel switches in a fabric if you:

- Use the firmware versions listed. All like switch models must use the same version when configured in a single-fabric or multi-fabric SAN. See Table 21 (page 85) through Table 23 (page 87). When updating switch firmware, you can use two successive switch firmware versions temporarily in a single-fabric and multi-fabric SAN.
- Follow the fabric rules. See “Fibre Channel switch fabric rules” (page 94).

For the latest information on supported B-series Fibre Channel switches and firmware versions, see the HP Storage Networking website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

**Table 21 B-series Fibre Channel switches**

HP switch name	Firmware versions	Fabric Management		Number of ports	
		HP Network Advisor	DCFM <sup>1</sup>		
HP SN8000B 8-Slot SAN Backbone Director Switch	7.0.0b 7.0.0a 6.4.2a 6.4.1b 6.4.1a 6.4.1 6.4.0c 6.3.2d 6.3.2b 6.3.2a 6.2.2e 6.2.2d	11.1.1 11.1.0	N/A	32 to 512	
HP SN8000B 4-Slot SAN Director Switch			N/A	32 to 256	
HP SN6000B 16Gb Fibre Channel Switch			N/A	24 to 48	
HP StorageWorks Encryption SAN Switch				32	
HP StorageWorks Encryption FC Blade				16	
Brocade 8Gb SAN Switch for HP BladeSystem c-Class				1 to 16 internal 1 to 8 external	
HP StorageWorks EVA4400 Embedded Switch Module, 8 Gb Brocade				10.4.5a 10.4.5 10.4.2 10.4.1 10.4.0 10.3.4 <sup>2</sup>	2 internal 10 external
HP StorageWorks DC SAN Backbone Director					16 to 512
HP StorageWorks 8/8 SAN Switch					8 to 24
HP StorageWorks 8/24 SAN Switch					16 to 24
HP StorageWorks 8/40 SAN Switch					24 to 40
HP StorageWorks 8/80 SAN Switch					48 to 80
HP StorageWorks DC04 SAN Director					16 to 256

<sup>1</sup> DCFM = Data Center Fabric Manager

<sup>2</sup> The Encryption SAN Switch and Encryption FC Blade require DCFM Enterprise V10.1.2 (or later) or DCFM Professional, Professional Plus V10.3.0 (or later).

**Table 22 B-series Fibre Channel switches and routers**

HP switch name	Firmware versions	Fabric Management			Number of ports
		FM <sup>1</sup>	HP Network Advisor	DCFM <sup>2</sup>	
HP StorageWorks 1606 Extension SAN Switch	7.0.0b 6.4.1 7.0.0a 6.4.0c	N/A	11.1.1 11.1.0	10.4.5a 10.4.5 10.4.2 10.4.1 10.4.0 10.3.4 <sup>3</sup>	16 at 8 Gb/s Fibre Channel 6 at 1-GbE FCIP
HP Multi Protocol Extension Blade	6.4.2a 6.3.2d 6.4.1b 6.3.2b 6.4.1a 6.3.2a				12 at 8 Gb/s Fibre Channel 10 at 1-GbE FCIP 2 at 10-GbE FCIP
HP Multi-protocol Router Blade for B-Series	6.4.2a 6.3.2d 6.4.1b 6.3.2b 6.4.1a 6.3.2a 6.4.1 6.2.2e 6.4.0c 6.2.2d	5.5 <sup>4</sup>			16 at 1, 2, or 4 Gb/s Fibre Channel 2 at 1-GbE FCIP

<sup>1</sup> FM = Fabric Manager

<sup>2</sup> DCFM = Data Center Fabric Manager

<sup>3</sup> The Encryption SAN Switch and Encryption FC Blade require DCFM Enterprise V10.1.2 (or later) or DCFM Professional, Professional Plus V10.3.0 (or later).

<sup>4</sup> Using HP StorageWorks Fabric Manager (FM) 5.5 or earlier with Fabric OS 6.1.1 (or later) can result in problems with Zoning, Performance monitoring, SYSLOG configuration, FICON wizards, IOD Delay, AAA replication, Port enable/disable, and other functions. Fabric Manager is not recommended for fabrics running these versions of Fabric OS, and is not supported with Fabric OS 6.3 (or later).

**Table 23 B-series legacy Fibre Channel switches and routers**

Legacy HP switch name	Firmware versions	Fabric Management			Number of ports
		FM <sup>1</sup>	HP Network Advisor	DCFM <sup>2</sup>	
HP StorageWorks iSCSI Director Blade					8 at 1, 2, or 4 Gb/s Fibre Channel 8 at 1-GbE iSCSI
HP StorageWorks 4/256 SAN Director base/power pack					16 to 384
HP StorageWorks 400 Multi-protocol Router	6.4.2a 6.3.2d 6.4.1b 6.3.2b				16 at 1, 2, or 4 Gb/s Fibre Channel 2 at 1-GbE
HP StorageWorks 4/32B SAN Switch full/power pack	6.4.1a 6.3.2a 6.4.1 6.2.2e				32
HP StorageWorks 4/32B SAN Switch base	6.4.0c 6.2.2d			10.4.5a 10.4.5	16 to 32
HP StorageWorks SAN Switch 4/32 base/full			11.1.1 11.1.0	10.4.2 10.4.1	16/32
HP StorageWorks 4/32 SAN Switch power pack				10.4.0 10.3.4 <sup>4</sup>	32
HP StorageWorks 4/64 SAN Switch base/power pack		5.5 <sup>3</sup>			32/48/64
HP StorageWorks 4/64 SAN Switch full					64
HP StorageWorks 4/8 SAN Switch base/full					8
HP StorageWorks 4/16 SAN Switch full/power pack	6.2.2e				16
Brocade 4Gb SAN Switch for HP c-Class BladeSystem	6.2.2d				1 to 16 internal 1 to 8 external
Brocade 4Gb SAN Switch for HP p-Class BladeSystem					8 internal 4 external
HP StorageWorks MSA SAN Switch 2/8	3.2.1c <sup>5</sup>		N/A	N/A	8
HP StorageWorks Multi-protocol Router <sup>6,7</sup>	7.4.1f 7.4.1e		N/A	N/A	8 or 16 at 1 or 2 Gb/s Fibre Channel or 1-GbE See <a href="#">Table 36 (page 110)</a> for scalability rules.

<sup>1</sup> FM = Fabric Manager

<sup>2</sup> DCFM = Data Center Fabric Manager

<sup>3</sup> Using HP StorageWorks Fabric Manager (FM) 5.5 or earlier with Fabric OS 6.1.1 (or later) can result in problems with Zoning, Performance monitoring, SYSLOG configuration, FICON wizards, IOD Delay, AAA replication, Port enable/disable, and other functions. Fabric Manager is not recommended for fabrics running these versions of Fabric OS, and is not supported with Fabric OS 6.3 (or later).

<sup>4</sup> The Encryption SAN Switch and Encryption FC Blade require DCFM Enterprise V10.1.2 (or later) or DCFM Professional, Professional Plus V10.3.0 (or later).

<sup>5</sup> You cannot manage these switches with DCFM.

<sup>6</sup> A minimum of firmware 7.4.1 is required for fabric compatibility with switch firmware 5.1x (or later).

<sup>7</sup> You cannot manage this router with DCFM.

## Features

Features of the B-series Fibre Channel switches include:

- **Access Gateway Mode**—Provides connectivity to B-series or other switch-series fabrics that support NPIV. The following B-series Fibre Channel switches support Access Gateway Mode: Brocade 8Gb SAN Switch for HP BladeSystem c-Class, HP StorageWorks SAN Switch 8/24, 8/40 (with 40 ports enabled and 6.2.0a or later), HP StorageWorks 2408 FCoE Converged Network Switch (Fibre Channel ports with 6.3.1a or later), and HP SN6000B 16Gb FC Switch.
- **Advanced Performance Monitor**—Analyzes resource utilization throughout the fabric.
- **Advanced WebTools**—Centralizes and simplifies switch management through a browser-based application.
- **Advanced Zoning**—Provides secure access control over fabric resources. Uses the switch hardware to enforce port and WWN zoning.
- **Extended Fabrics**—Enhances performance for Fibre Channel SAN connectivity for distances greater than 1 km (8 Gb/s), 2.5 km (4 Gb/s), 5 km (2 Gb/s), and 10 km (1 Gb/s) to improve disaster recovery operations and ensure business continuity. Extended fabric support for distances up to 10 km is provided in all B-series switches. For distances greater than 10 km, an Extended Fabrics license is required. For maximum supported distances based on the link speed, see “SAN fabric connectivity and switch interoperability rules” (page 146).
- **Fabric Manager**—Centralizes fabric management through a host-based application. This is a legacy product, currently in sustaining mode only.
- **Data Center Fabric Manager (DCFM)**—Centralizes fabric management through a host-based application. This is available free for single-fabric management (DCFM Professional), or for purchase for multiple-fabric management and increased device count (DCFM Professional Plus and DCFM Enterprise).

---

**NOTE:** See [Table 21 \(page 85\)](#) through [Table 23 \(page 87\)](#) for switch models supported with Fabric Manager, DCFM, and HP Network Advisor.

---

- **HP Network Advisor**—Centralizes fabric management through a host-based application. This is available free for single-fabric management (Network Advisor Professional), or for purchase for multiple-fabric management and increased device count (Network Advisor Professional Plus and Network Advisor Enterprise).
- **Fabric Watch**—Proactively monitors the health and performance of switches and the fabric.
- **ISL Trunking**—Combines multiple links between switches to form a single, logical ISL with a total bandwidth of 64 Gb/s. This feature enables dynamic load balancing of data across ISLs.
- **FCIP Trunking**—Allows multiple IP source and destination address pairs (defined as FCIP circuits) via multiples of the 1-GbE or 10-GbE interfaces to provide a high-bandwidth FCIP tunnel and failover resiliency. FCIP tunnels using 1-GbE ports can have up to four FCIP circuits spread over four 1-GbE ports. FCIP tunnels using 10-GbE ports can have up to ten FCIP circuits per 10-GbE port. Each FCIP circuit supports four QoS classes as a TCP connection (Class-F, Hi, Medium, and Low priority).



---

**NOTE:** FCIP Trunking is supported only on the 1606 Extension SAN Switch and the DC Dir Switch MP Extension Blade. For more information, see [“B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade”](#) (page 280).

---

- **Adaptive Rate Limiting**—Dynamically adjusts data transmission through the TCP connections in a FCIP tunnel, based on preconfigured minimum and maximum data rates and by taking feedback from the quality of the TCP connections. This allows FCIP connections to use the maximum bandwidth available while providing a minimum bandwidth guarantee.

---

**NOTE:** ARL is supported for configuring FCIP tunnels only on the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade.

---

- **NPIV**—See [“SAN fabric connectivity and switch interoperability rules”](#) (page 146). The following B-series Fibre Channel switches are NPIV compliant:
  - **16 Gb/s:** SN8000B 8-Slot SAN Backbone Director, SN8000B 4-Slot SAN Director, HP SN6000B 16Gb Fibre Channel Switch
  - **8 Gb/s:** Brocade 8Gb SAN Switch for HP BladeSystem c-Class; SAN Switch 8/8, 8/24, 8/40, and 8/80; DC SAN Backbone Director; DC04 SAN Director; EVA4400 Embedded Switch Module; Encryption SAN Switch; Encryption FC Blade; 2408 FCoE Converged Network Switch (fibre channel ports)
  - **4 Gb/s:** 4/8, 4/16, 4/32, 4/32B, 4/64, and 4/256, Brocade 4Gb SAN Switch for HP c-Class BladeSystem; and Brocade 4Gb SAN Switch for HP p-Class BladeSystem
  - **2 Gb/s:** 2/128
- **Secure Fabric OS**—Provides flexible security and policy-based administration to protect data from unauthorized access and corruption. All Secure Fabric OS features are included in the base Fabric OS beginning with 5.3.0 firmware. Secure Fabric OS can coexist in a fabric running Fabric OS 5.3.x firmware, but this is not supported with later versions of Fabric OS.

Features of Fabric OS 6.x follow (see the appropriate Fabric OS administrator guide for configuration restrictions and more information):

- **Adaptive Networking with QoS (optional licensed feature)**—Provides tools to support high bandwidth on high-priority connections, ensuring optimum performance, even in congested networks. This option includes two features: QoS SID/DID Prioritization and Ingress Rate Limiting. Fabric OS 6.3 includes QoS support for domain, index and WWN zone definitions, and bottleneck detection for F\_Ports. For configuration restrictions, see the *HP StorageWorks Fabric OS 6.x Administrator Guide*.
- **ICLs (DC SAN Backbone Directors and DC04 SAN Directors only)**—Provides dedicated high-bandwidth links between two DC SAN Backbone Director chassis, without using the front-end 8 Gb/s ports. Each DC SAN Backbone Director must have the ICL license installed to enable ICL connections.
- **Top talkers (requires Advanced Performance Monitor license)**—Provides real-time information about the top bandwidth-consuming flows passing through a specific point in the network.
- **Traffic isolation zones**—Enables application segregation, guaranteeing separate paths through the fabric. You can isolate traffic assigned to ISLs in the fabric.
- **Virtual Fabrics**—Enables division of a single physical chassis into logical switches, allowing you to create logical fabrics with other switches. This feature was introduced with Fabric OS 6.2x and is supported on the DC SAN Backbone Director, DC04 SAN Director, and 8/40 and 8/80 SAN switches. Virtual Fabrics is enabled by default on these switches with V6.3.1a (or later) factory installed.

- **Fibre Channel routing and FCIP enhancements**—Support for EX\_Port Masterless Trunking and QoS, enhanced compression, FCIP Trunking, and ARL features for the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade.

---

**NOTE:** Masterless EX\_Port Trunking is not supported when using FCIP links in Fabric OS 6.3.0. It is supported in Fabric OS 6.3.0a (or later).

---

- **Encryption enhancements**—Supports the System card and Quorum card, tape encryption and HA clustered key vault beginning with Fabric OS 6.3.0.

---

**NOTE:** Advanced Zoning and WebTools are automatically enabled without additional licensing starting with Fabric OS 6.1x.

---

Features of Fabric OS 7.x follow (see the *HP B-series Fabric OS 7.x Release Notes* and appropriate Fabric OS administrator guide for configuration restrictions and more information):

- **In-flight encryption and compression on 16 Gb/s capable ISLs**—Provides security for frames while they are passing between two switches. Compression improves bandwidth utilization on the ISLs, especially over long distance. Encryption and compression are both disabled by default but may be enabled at an individual port level.
- **Buffer credit loss detection and automatic recovery on 16 Gb/s capable ISLs**—Provides support for the detection and automatic recovery of buffer credits on 16 Gb/s capable ISLs (both ends of the link must terminate on 16 Gb/s capable ports).
- **10 Gb/s Fibre Channel capability on 16 Gb/s Fibre Channel platforms**—Provides the ability to configure ports on the SN8000B 16 Gb/s 32-port or 48-port Fibre Channel Blades for the HP SN8000B 8-Slot SAN Backbone Director Switch, the HP SN8000B 4-Slot SAN Director Switch, and the HP SN6000B 16Gb FC Switch to operate at 10 Gb/s.
- **Advanced diagnostics support on 16 Gb/s Fibre Channel platforms**—Provides new and enhanced diagnostics support, such as Diagnostic port (D\_Port), Forward Error Correction (FEC), real-time power monitoring, and advanced SFP monitoring.
- **Dynamic fabric provisioning**—Supports Fabric Assigned WWN (FA-WWN), which uses a virtual WWN for a server instead of the server's physical pWWN to create zoning and LUN mapping/masking (requires specific Brocade HBA and driver support).
- **Detection and resolution of duplicate device WWNs**—Implements a feature to detect duplicate device WWNs in a fabric.
- **FCIP enhancements**—Some of the FCIP enhancements include:
  - 10 GigE lossless failover on HP DC SAN Director Multiprotocol Extension Blade
  - Multi-Gigabit circuits on HP DC SAN Director Multiprotocol Extension Blade, which allows more than 1 Gb/s to be configured on a single FCIP Circuit
  - 10 GigE Adaptive Rate Limiting (ARL) on HP DC SAN Director Multiprotocol Extension Blade
- **Configurable QoS on FCIP tunnel**—Provides the option to override the default percentages of bandwidth assigned to each of the QoS classes within an FCIP tunnel.
- **Auto-mode option for compression**—Provides the option to automatically adjust the compression mode (1, 2, or 3) based on the maximum configured tunnel bandwidth on HP DC SAN Director Multiprotocol Extension Blade and HP 1606 SAN Extension Switch.
- **XISL support on HP DC SAN Director Multiprotocol Extension Blade**—Provides the ability to use VE ports as XISLs on HP DC SAN Director Multiprotocol Extension Blade only in both 1 Gb/s and 10 Gb/s modes.

- **InBand Management on HP DC SAN Director Multiprotocol Extension Blade and HP 1606 SAN Extension Switch**—Allows you to define one or more IP addresses on the non-management GE ports and provide a management path from WAN to the switch CP.
- **Encryption platform enhancements (HP StorageWorks Encryption SAN Switch/HP StorageWorks DC Switch Encryption FC Blade)**—Fabric OS 7.0.0a (or later) include support for the key vault diagnostics. The key vault connectivity information is periodically collected and any connectivity errors are reported as RASlog.
- **Security enhancements**—Provides the following security enhancements:
  - SSH authentication using public keys
  - SFTP support for firmwaredownload, configupload and configdownload
  - IPv6 support for LDAP authentication
- **FCoE enhancements**—Provides support for high availability for FCoE traffic going through HP DC SAN Director 10/24 FCoE Blade in HP StorageWorks DC SAN Backbone Director Switch/HP StorageWorks DC04 SAN Director Switch.

Table 24 (page 91) provides a comparison of the high-availability features for B-series Fibre Channel switches.

**Table 24 B-series Fibre Channel switch high-availability feature comparison**

Model	Redundant /hot-swappable power	Redundant /hot-swappable cooling	Redundant core switching blade	Redundant control processor	Non-disruptive code activation	Non-disruptive port expansion	Redundant active components
HP SN8000B 8-Slot SAN Backbone Director Switch	Yes/Yes	Yes/Yes	Yes	Yes	Yes	Yes	Yes
HP SN8000B 4-Slot SAN Director Switch	Yes/Yes	Yes/Yes	Yes	Yes	Yes	Yes	Yes
HP SN6000B 16Gb Fibre Channel Switch	Yes/Yes	Yes/Yes	No	No	Yes	Yes	No
Encryption SAN Switch	Yes/Yes	Yes/Yes	No	No	Yes	N/A	No
Brocade 8Gb SAN Switch for HP BladeSystem c-Class	N/A	N/A	No	No	Yes	Yes	No
Brocade 4Gb SAN Switch for HP c-Class BladeSystem	N/A	N/A	No	No	Yes	Yes	No
Brocade 4Gb SAN Switch for HP p-Class BladeSystem	N/A	N/A	No	No	Yes	No	No
SAN Switch 4/8, 4/16	No/No	Yes/No	No	No	Yes	Yes	No
SAN Switch 4/32B, 4/32, 4/64	Yes/Yes	Yes/Yes	No	No	Yes	Yes	No
EVA4400 Embedded Switch Module, 8 Gb Brocade	N/A	N/A	No	No	Yes	No	Yes (2 switches per chassis)
SAN Switch 8/8, 8/24, 8/40, 8/80	Yes/Yes	Yes/Yes	No	No	Yes	Yes	No
DC SAN Backbone Director	Yes/Yes	Yes/Yes	Yes	Yes	Yes	Yes	Yes
DC04 SAN Director	Yes/Yes	Yes/Yes	Yes	Yes	Yes	Yes	Yes

**Table 24 B-series Fibre Channel switch high-availability feature comparison** *(continued)*

<b>Model</b>	<b>Redundant /hot-swappable power</b>	<b>Redundant /hot-swappable cooling</b>	<b>Redundant core switching blade</b>	<b>Redundant control processor</b>	<b>Non-disruptive code activation</b>	<b>Non-disruptive port expansion</b>	<b>Redundant active components</b>
SAN Director 2/128, 4/256	Yes/Yes	Yes/Yes	No	Yes	Yes	Yes	Yes
400 MP Router	Yes/Yes	Yes/Yes	No	No	Yes	No	No
MP Router	Yes/Yes	Yes/Yes	No	No	No	Yes	No
HP StorageWorks 1606 Extension SAN Switch	Yes/Yes	Yes/Yes	No	No	Yes	No	No

## Usage

Table 25 (page 93) describes the use of B-series Fibre Channel switches as core switches.

**Table 25 Using B-series Fibre Channel switches as core switches**

Model	1-96 user ports	97-224 user ports	225-500 user ports	501-728 user ports	729-1280 user ports
HP SN8000B 8-Slot SAN Backbone Director Switch HP SN8000B 4-Slot SAN Director Switch DC SAN Backbone Director DC04 SAN Director Director Switch 2/128, 4/256	Excellent				
SAN Switch 8/80, 4/64	Excellent		Very good	Good	Not recommended
HP SN6000B 16Gb FC Switch SAN Switch 8/24, 8/40, 4/32B, 4/32	Excellent	Very good	Good	Not recommended	
SAN Switch 8/8, 400 MP Router, 1606 Extension SAN Switch	Good			Not recommended	
Brocade 8Gb SAN Switch for HP BladeSystem c-Class Brocade 4Gb SAN Switch for HP p-Class BladeSystem Brocade 4Gb SAN Switch for HP c-Class BladeSystem EVA4400 Embedded Switch Module, 8 Gb Brocade	Good	Not recommended			

Table 26 (page 94) describes the use of B-series Fibre Channel switches as edge switches.

**Table 26 Using B-series Fibre Channel switches as edge switches**

Model	1-96 user ports	97-224 user ports	225-500 user ports	501-728 user ports	729-1280 user ports
HP SN8000B 8-Slot SAN Backbone Director Switch HP SN8000B 4-Slot SAN Director Switch DC SAN Backbone Director DC04 SAN Director SAN Director 4/256 Core Switch 2/128 Director Switch 2/128, 4/256			Excellent		
HP SN6000B 16Gb FC Switch SAN Switch 8/40, 8/80, 4/32B, 4/32, 4/64			Excellent		
SAN Switch 8/8, 8/24, 4/8, 4/16 Brocade 8Gb SAN Switch for HP BladeSystem c-Class Brocade 4Gb SAN Switch for HP p-Class BladeSystem Brocade 4Gb SAN Switch for HP c-Class BladeSystem EVA4400 Embedded Switch Module, 8 Gb Brocade 400 MP Router 1606 Extension SAN Switch	Excellent		Very good		Good

## Fibre Channel switch fabric rules

This section describes the fabric rules for B-series Fibre Channel switches and routers, and other factors you should consider when building B-series fabrics.

When using B-series Fibre Channel switches in a fabric:

- Use the HP default settings for all current HP switches.
- Use the legacy HP default settings if the fabric contains only legacy HP switches.
- Use the legacy Compaq default settings if the fabric contains only legacy Compaq switches or a mix of legacy Compaq and legacy HP switches.

---

**NOTE:** Contact HP Services for the configuration files that contain the appropriate switch settings.

---

This section describes the following topics:

- “Operating systems and storage products” (page 95)
- “B-series Fibre Channel switch and fabric rules” (page 95)
- “B-series Encryption Switch fabric rules” (page 99)
- “1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules” (page 100)
- “Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules” (page 103)

- “MP Router fabric rules” (page 107)
- “Core switch addressing mode” (page 112)
- “Zoning limits and enforcement” (page 112)
- “Zoning guidelines for B-series Fibre Channel switches” (page 113)

## Operating systems and storage products

The fabric rules for B-series switches and MP Routers apply to SANs that include the operating systems and storage products listed in [Table 27 \(page 95\)](#).

**Table 27 B-series Fibre Channel switch operating system and storage system support**

Operating systems	Storage products
<ul style="list-style-type: none"> <li>• HP-UX</li> <li>• OpenVMS</li> <li>• Tru64 UNIX</li> <li>• Apple Mac OS X</li> <li>• Citrix Xen</li> <li>• IBM AIX</li> <li>• Linux</li> <li>• Microsoft Windows</li> <li>• Sun Solaris</li> <li>• VMware ESX</li> </ul>	<ul style="list-style-type: none"> <li>• EVA8400/6400</li> <li>• EVA4400/4400 (embedded switch)</li> <li>• EVA8100/8000/6100/6000/4100/4000</li> <li>• EVA5000/3000</li> <li>• MSA2000fc G2 (MSA2300fc)/MSA2000fc</li> <li>• MSA1500/1000</li> <li>• P6300/P6500 EVA</li> <li>• P9500</li> <li>• P2000 G3 FC</li> <li>• VA7410/7400/7110/7100</li> <li>• XP24000/20000</li> <li>• XP12000/10000</li> <li>• XP1024/128</li> <li>• XP512/48</li> <li>• 3PAR F-Class, T-Class</li> <li>• HP P9000 (XP) Continuous Access</li> <li>• HP P6000 Continuous Access</li> </ul> <p>See “<a href="#">HP P6000 Continuous Access SAN integration</a>” (page 220)</p>

The operating systems listed in [Table 27 \(page 95\)](#) may not be supported with all the storage systems listed. For current operating system and storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

See “[Heterogeneous server rules](#)” (page 159) and the following storage system rules chapters:

- “[P2000 and MSA storage system rules](#)” (page 203)
- “[P6000/EVA storage system rules](#)” (page 213)
- “[P9000/XP and VA storage system rules](#)” (page 227)

## B-series Fibre Channel switch and fabric rules

The following fabric rules apply to all B-series Fibre Channel SANs. They also apply, in general, to HP P9000 (XP) Continuous Access and HP P6000 Continuous Access (formerly HP StorageWorks Continuous Access EVA) configurations. However, additional rules apply to HP P6000 Continuous Access implementations. For more information, see the *HP P6000 Continuous Access Implementation Guide*.

[Table 28 \(page 96\)](#) describes the rules for creating a Fibre Channel SAN with B-series Fibre Channel switches.

**Table 28 B-series Fibre Channel switch fabric rules**

Rule number	Description
1	<p>A maximum of 56 switches and 2,560 virtual or physical connections in a single fabric with HP SN8000B 8-Slot SAN Backbone Director Switch, HP SN8000B 4-Slot SAN Director Switch, or HP SN6000B 16Gb FC Switch using firmware 7.x (or later) or DC SAN Backbone, or DC04 SAN Directors, 8/8, 8/24, 8/40, 8/80, 4/8, 4/16, 4/32B, 4/32, 4/64, or 4/256 switches using firmware 6.x (or later), or 2/64 or 2/128 switches using firmware 5.3.x (or later).</p> <p>A maximum of 16 DC SAN Backbone Director or DC04 SAN Director chassis, or 6,000 virtual or physical connections (WWNs logged into a single fabric), are supported in a fabric containing HP SN8000B 8-Slot Backbone Director, HP SN8000B 4-slot Backbone Director, DC SAN Backbone Directors or DC04 SAN Directors only.</p> <p>A maximum of 4,096 virtual or physical connections are supported in a fabric containing other switch models using firmware 6.1.x (or later).</p>
2	<p>A maximum of 56 switches and 2,560 virtual or physical connections in a single fabric with 4/8, 4/16, 4/32B, 4/32, 4/64, 4/256, 2/64, or 2/128 switches using firmware 5.x (or later).</p>
3	<p>8 Gb/s switches and ARB(FF)/ARB(FF) mode setting—HP recommends all 8 Gb/s switches running FOS 6.3.1a and later (6.4.0a and later for XP storage array connectivity) be set to mode 3 using the <code>portcfgfillword</code> command. For more information, see the <i>Fabric OS Command Reference Manual supporting Fabric OS v6.4</i> and the FOS release notes.</p> <p><b>Note:</b> HP SN8000B 8-Slot SAN Backbone Director Switch, HP SN8000B 4-Slot SAN Director Switch, or the HP SN6000B 16Gb FC Switch automatically selects the proper fill word.</p>
4	<p>HP StorageWorks 4/256 SAN Director—Maximum of 20 chassis per fabric with all 16-port Fibre Channel blades, maximum of 10 chassis with all 32-port Fibre Channel blades. For all configurations, the maximum number of total virtual or physical connections must not exceed 2,560. Each chassis adds one switch to the fabric switch count if it is configured as one domain. This switch does not support port-based routing. See "Routing Traffic" in the <i>HP StorageWorks Fabric OS 5.x Administrator's Guide</i>.</p>
5	<p>HP StorageWorks 2/128 SAN Director—Maximum of 20 chassis per fabric. Each chassis adds one switch to the fabric switch count if it is configured as one domain.</p>



**Table 28 B-series Fibre Channel switch fabric rules** (continued)

Rule number	Description
6	<p>B-series switches and routers supported routing modes with Continuous Access products:</p> <ul style="list-style-type: none"> <li>• HP P9000 (XP) Continuous Access—Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3). See “<a href="#">HP P9000 (XP) Continuous Access FCIP gateway support</a>” (page 307).</li> <li>• HP P6000 Continuous Access—Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later). See “<a href="#">Network requirements for long-distance IP gateways with XCS 10x or XCS 09x</a>” (page 289).</li> </ul> <p>Routing policy—HP recommends that you use the same routing policy in all switches and routers within a fabric.</p>
7	<p>EVA4400 Embedded Switch Module, 8 Gb Brocade—Supports connectivity to:</p> <ul style="list-style-type: none"> <li>• Servers</li> <li>• Storage</li> </ul> <p><b>Note:</b> The EVA4400 Embedded Switch Module, 8 Gb Brocade supports the same connectivity to HP storage systems as the 8/8, 8/24, and 8/80 switches.</p> <ul style="list-style-type: none"> <li>• Other B-series switches, including the EVA4400 Embedded Switch Module, 8 Gb Brocade</li> <li>• MPX200 Multifunction Router with iSCSI (see “<a href="#">MPX200 Multifunction Router with iSCSI for EVA storage</a>” (page 324)) and EVA and EVA4400 iSCSI Connectivity Option (see “<a href="#">EVA and EVA4400 iSCSI Connectivity Option</a>” (page 333))</li> <li>• MPX200 Multifunction Router with FCIP, HP IP Distance Gateway, and B-series 400 MP Router (see “<a href="#">HP multi-protocol long-distance products</a>” (page 265))</li> </ul>
8	<p>Access Gateway (AG) requirements are:</p> <ul style="list-style-type: none"> <li>• AG mode for Brocade 8Gb SAN Switch for HP BladeSystem c-Class and HP StorageWorks SAN Switch 8/8 or 8/24 (with 24 ports enabled) requires firmware 6.1x (or later).</li> <li>• AG mode for the SAN Switch 8/40 requires firmware 6.2x (or later).</li> <li>• AG mode for the 2408 FCoE Converged Network Switch requires firmware 6.3.1a (or later).</li> <li>• AG mode for Brocade 4Gb SAN Switch for HP c-Class and p-Class BladeSystems requires firmware 5.2.1b (or later).</li> <li>• AG mode for HP StorageWorks SAN Switch 4/8 or 4/16 requires firmware 5.3.0 (or later).</li> <li>• All switches that communicate with Access Gateway switches require firmware 5.1.x (or later). See “<a href="#">BladeSystem with Brocade Access Gateway mode</a>” (page 163).</li> </ul>
9	<p>In a fabric that contains a Brocade 4Gb SAN Switch for HP c-Class or p-Class BladeSystem, StorageWorks SAN Switch 4/8, 4/16, 4/32B, 4/32, 4/64, SAN Director 2/128, or SAN Director 4/256, Core switch addressing mode is required on all other switches in that fabric. See “<a href="#">Core switch addressing mode</a>” (page 112).</p>
10	<p>All B-series Fibre Channel switches and routers must use the HP default setting of <code>dlsset</code> for non-Continuous Access configurations and HP P6000 Continuous Access configurations running XCS version 09534000 (or later) have two protocol options—HP Fibre Channel protocol or HP SCSI Fibre Channel protocol. The settings for DLS and IOD functionality on the B-series switches should be configured as follows:</p> <p>HP Fibre Channel protocol:</p> <ol style="list-style-type: none"> <li>1. <code>dlsreset</code> (DLS disabled)</li> <li>2. <code>iodset</code> (IOD enabled)</li> </ol> <p>HP SCSI Fibre Channel protocol:</p> <ol style="list-style-type: none"> <li>1. <code>dlsset</code> (DLS enabled)</li> <li>2. <code>iodreset</code> (IOD disabled)</li> </ol> <p>Note: <code>dlsset</code> is the default policy when <code>aptpolicy 3</code> (enable exchange-based routing) is enabled on B-series switches.</p> <p>For more information, see the <i>HP P6000 Continuous Access Implementation Guide</i>.</p> <p>For HP P9000 (XP) Continuous Access, see “<a href="#">HP P9000 (XP) Continuous Access FCIP gateway support</a>” (page 307).</p>

**Table 28 B-series Fibre Channel switch fabric rules** (continued)

Rule number	Description
11	<p>Within a fabric, assign a unique domain number (domain ID) and a unique WWN to each switch. All switch configuration parameters for the same switch models must be the same. Do not configure any switches with a domain ID of 8, which is reserved for HP-UX.</p> <p>Insistent Domain ID mode is required on all B-series switches in fabrics that include HP-UX 11i v3 (or earlier) or IBM AIX servers, when configuring switches for FICON or when implementing port-based fabric zoning.</p> <p>Routing policy—HP recommends the same routing policy be used in all switches within a fabric.</p>
12	<p>Switch data routing policies—Different switch models support different routing policies. For most configurations, the default routing policy provides the best performance. See "Routing Traffic" in the appropriate Fabric OS administrator guide.</p>
13	<p>Maximum of seven hops (eight switches) between any two communicating devices.</p>
14	<p>You can configure a Fibre Channel port as an F_Port, FL_Port, E_Port, or EX_Port (Fibre Channel routing), and configure a GbE port as a VE_Port (FCIP) or VEX_Port (FCIP with Fibre Channel routing).</p> <p><b>Note:</b> The 16 Gb/s port blades do not support FL ports.</p> <p>Traffic switching is supported between the following port types:</p> <ul style="list-style-type: none"> <li>• Fibre Channel switching (F_Port, FL_Port, E_Port)</li> <li>• Fibre Channel switching and Fibre Channel routing (F_Port, FL_Port, EX_Port)</li> <li>• Fibre Channel switching and FCIP (F_Port, FL_Port, VE_Port, VEX_Port)</li> <li>• Fibre Channel routing and FCIP (VE_Port, VEX_Port)</li> </ul> <p><b>Note:</b> VEX_Ports and EX_Ports cannot be connected to the same edge fabric.</p>

**NOTE:** Not all topologies can support the maximum port or switch count. See “Topology maximums” (page 33).

### Switch database size

Table 29 (page 98) describes the database size rules for B-series Fibre Channel switches in a fabric.

**Table 29 Database size rules for B-series Fibre Channel switches**

Rule number	Description
1	<p><b>Secure Fabric OS</b>—With security enabled in a fabric, the maximum security database size is:</p> <ul style="list-style-type: none"> <li>• 128 KB with 64 KB active if the fabric contains only 2 Gb/s or 4 Gb/s switches.</li> <li>• 1 MB with Fabric OS 6.x.</li> </ul> <p>This database, which stores DCC, SCC, and FCS policies, can support ACL security settings for environments with up to 14,000 attached devices. This is not supported with the 4/16 SAN switch.</p>
2	<p><b>Advanced Zoning</b>—The maximum zoning database size is as follows:</p> <ul style="list-style-type: none"> <li>• If all switches in a fabric are using firmware 5.2x (or later), the database size must not exceed 1 MB.</li> <li>• If the fabric contains 2 Gb/s or 4 Gb/s switches with firmware 3.2.x (or later), the database must not exceed 256 KB.</li> </ul> <p><b>Note:</b> Use the <code>cfgsize</code> command to determine the size of the zoning database.</p>

### ISL maximums

You can use all ports on all B-series Fibre Channel switches for ISLs, with a maximum of one half of the total ISL port count configured to the same destination switch.

**NOTE:** Some switches require licensing for additional ISL ports.

## DC SAN Director FC8-64 cabling

The DC SAN Director FC8-64 port blade requires mSFP fiber optic cable connectors. Table 30 (page 99) describes the two possible connection scenarios and solutions for cabling.

**Table 30 DC SAN Director High Density 64-port blade mSFP cabling**

Scenario	Solution
Connecting from an FC8-64 port blade (mSFP) to an existing standard Fibre Channel cabling infrastructure (LC)	Use an HP mSFP-LC fiber optic cable connected directly to the LC port of an existing LC-based switch or device to an HP LC-LC coupler then connected to an existing LC fiber optic cable, or to an LC patch panel.
Connecting from an FC8-64 port blade (mSFP) to another FC8-64 (mSFP) for ISL connectivity	<ul style="list-style-type: none"> <li>Option 1: Use two mSFP-LC fiber optic cables connected together with an LC-LC coupler.</li> <li>Option 2 (Brocade): Use an mSFP-mSFP fiber optic cable.</li> </ul> <p><b>Note:</b> The mSFP-mSFP cable is not available from HP. See the HP DC SAN Director QuickSpec and Brocade <i>Best Practices Guide: High Density Cable Management Solutions</i>.</p>

## B-series Encryption Switch fabric rules

This section describes the fabric rules for the B-series Encryption Switches (HP StorageWorks Encryption SAN Switch and HP StorageWorks Encryption FC Blade) and other factors you should consider when building fabrics. These switches have a built-in encryption engine and contain 8 Gb/s Fibre Channel ports (32 for the encryption switch, 16 for the encryption blade) and two GbE ports.

Table 31 (page 99) describes the rules for creating fabrics with B-series Encryption Switches. For more information about the security uses of these switches, see “B-series Encryption Switch and Encryption FC Blade security” (page 366). HP strongly recommends that you read the *Brocade Fabric OS Encryption Administrator's Guide* at [http://h18006.www1.hp.com/storage/saninfrastructure/switches/encrypt\\_sanswitch/index.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/storage/saninfrastructure/switches/encrypt_sanswitch/index.html?jumpid=reg_R1002_USEN) before configuring a fabric using the encryption switch or encryption blade.

**Table 31 B-series Encryption Switch fabric rules**

Rule number	Description
1	Encryption blade can be installed in an HP SN8000B 8-Slot SAN Backbone Director Switch, HP SN8000B 4-Slot SAN Director Switch, DC SAN Director, or DC04 SAN Director.
2	There is a maximum of four encryption blades per director.
3	There is a maximum of four encryption switches per encryption group.
4	There is a maximum of 16 encryption engines (with four Directors, each having four encryption engines) per encryption group.
5	There is a maximum of 256 <sup>1</sup> physical targets per encryption switch or blade.
6	There is a maximum of 1,024 <sup>1</sup> physical initiators per encryption switch or blade.

<sup>1</sup> This value is half the value of the actual limits supported in a high-availability cluster.

## 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules

This section describes the fabric rules for the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade, and other factors you should consider when building B-series fabrics with Fibre Channel routing enabled. These fabric rules apply to SANs that include the same operating systems and storage products as the B-series switches. See [“Operating systems and storage products” \(page 95\)](#). For information about configuring the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade for FCIP, see [“B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade” \(page 280\)](#).

[Table 32 \(page 100\)](#) describes the rules for creating fabrics with Fibre Channel routing using the 1606 Extension SAN Switch and the DC Dir Switch MP Extension Blade.

**Table 32 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules**

Rule number	Description
1	The DC Dir Switch MP Extension Blade supports only the HP SN8000B 8-Slot SAN Backbone Director, HP SN8000B 4-Slot SAN Director, DC SAN Backbone Director and DC04 SAN Director. The maximum number of DC Dir Switch MP Extension Blades per HP SN8000B SAN Directors (both 4 and 8-slot) is four. The maximum number of DC Dir Switch MP Extension Blades per DC SAN Backbone Director or DC04 SAN Director is two.
2	All configurations must use the default setting for R_A_TOV (10,000 milliseconds) and E_D_TOV (2,000 milliseconds).
3	Devices connected to edge fabrics can be routed through an LSAN to a backbone fabric. <b>Note:</b> Edge fabric to backbone fabric routing is supported using the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, and MP Router Blade. It is not supported if an MP Router is configured in the backbone.
4	Direct connection of devices to the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade Fibre Channel ports (F_Port, FL_Port) is supported.
5	Fibre Channel server boot through the 1606 Extension SAN Switch or the DC Dir Switch MP Extension Blade is not supported.
6	Each 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade can attach to one backbone fabric only; multiple routers can attach to the same backbone fabric.
7	If a DC Dir Switch MP Extension Blade has a port configured as an EX_Port, then the blade Fibre Channel ports, the Fibre Channel ports on the director containing the blade, and all switches connected to the director using an E_Port are part of the backbone fabric.
8	A backbone fabric supports a maximum of 256 devices.
9	You can configure a Fibre Channel port as an F_Port, FL_Port, E_Port, or EX_Port (Fibre Channel routing), and configure a GbE port as a VE_Port (FCIP) or VEX_Port (FCIP with Fibre Channel routing). Traffic switching is supported between the following port types: <ul style="list-style-type: none"> <li>• F_Port, FL_Port, E_Port (Fibre Channel switching)</li> <li>• F_Port, FL_Port, EX_Port (Fibre Channel switching and Fibre Channel routing)</li> <li>• F_Port, FL_Port, VE_Port, VEX_Port (Fibre Channel switching and FCIP)</li> <li>• VE_Port, VEX_Port (Fibre Channel routing and FCIP)</li> </ul> <b>Note:</b> VEX_Ports and EX_Ports cannot be connected to the same edge fabric. VEX_Ports are not supported on the DC Dir MP Extension Blade with firmware 6.3.0.
10	FCIP is supported between pairs of 1606 Extension SAN Switches and DC Dir Switch MP Extension Blades, or between a 1606 Extension SAN Switch and a DC Dir Switch MP Extension Blade. FCIP is not supported in between the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade and the MP Router, 400 MP Router, or MP Router Blade. See <a href="#">“B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade” (page 280)</a> .

**Table 32 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade fabric rules** (continued)

Rule number	Description
11	FCIP tunnels (virtual E_Ports) pass traffic between the SAN extension devices. <ul style="list-style-type: none"> <li>• The 1606 Extension SAN Switch supports up to 8 FCIP tunnels.</li> <li>• The DC Dir Switch MP Extension Blade supports up to 20 FCIP tunnels.</li> </ul>
12	IPsec is supported on the 1606 Extension SAN Switch beginning with firmware 6.3.1a. See <a href="#">Rule 18</a> .
13	FCIP FastWrite and Fibre Channel FastWrite support: <ul style="list-style-type: none"> <li>• FCIP is supported with FCIP FastWrite SCSI write acceleration when using firmware 6.3.0a (or later) with HP P9000 (XP) Continuous Access. FCIP FastWrite is not supported with HP P6000 Continuous Access.</li> </ul> <p><b>Note:</b> FCIP FastWrite and IPsec are mutually exclusive and cannot be configured simultaneously.</p> <ul style="list-style-type: none"> <li>• Fibre Channel FastWrite—FC FastWrite is currently not supported with the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade.</li> </ul>
14	For redundancy, HP recommends using a minimum of two EX_Ports (router IFLs) for each fabric connected to a 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade. The number of EX_Ports required is based on the performance requirements for all devices shared by fabrics through the router. <ul style="list-style-type: none"> <li>• Calculate the number of EX_Ports using a 5:1 ratio of standard ISLs or switch device ports to router EX_Port IFLs. Monitor port performance to determine if this ratio is acceptable based on usage. Use the <code>portperfshow</code> command to monitor router port performance.</li> <li>• For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1.</li> </ul>
15	Supports routing between B-series secured and nonsecured edge fabrics using advanced security features or ACLs.
16	Supports routing between B-series and M-series/McDATA edge fabrics. See the <i>HP StorageWorks Fabric Interoperability: Merging Fabrics Based on M-series and B-series Fibre Channel Switches Application Notes</i> at <a href="http://hp.com/go/sdgmanuals">http://hp.com/go/sdgmanuals</a> . <p><b>Note:</b> Continuous Access products are not supported in routed fabrics using both B-series and M-series/McDATA edge fabrics.</p>
17	Supports the following features: <ul style="list-style-type: none"> <li>• 8 Gb/s on FC ports</li> <li>• FCIP Trunking</li> <li>• Adaptive Rate Limiting</li> <li>• 10-GbE ports (DC Dir Switch MP Extension Blade)</li> </ul>
18	The 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade do not support the following features with firmware 6.3.0: <ul style="list-style-type: none"> <li>• IPsec</li> <li>• DiffServ priorities</li> <li>• VLAN tagging</li> <li>• Third-party WAN optimization</li> <li>• IPv6 addresses for FCIP tunnels</li> <li>• VEX_Ports (DC Dir Switch MP Extension Blade)</li> </ul>

### Scalability rules

Table 34 (page 105) lists the scalability rules for Meta SANs using integrated Fibre Channel routing, the 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, and MP Router Blade. For MP Router scalability rules, see “Scalability rules” (page 110). For configurations

with an MP Router and a 400 MP Router or MP Router Blade, apply the rules listed in [Table 34 \(page 105\)](#).

The following terms describe 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade scalability:

- **Front phantom domains**—All EX\_Ports on a Fibre Channel router that are connected to the same edge fabric are presented as a single front phantom domain.
- **Translate phantom domains**—Assigned to each remote fabric connected to a 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, or MP Router Blade. Each fabric has one unique translate domain, regardless of how many EX\_Ports are connected to the fabric.
- **Backbone fabric**—The 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade backbone fabric consists of at least one 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade, and includes all MP Routers and B-series switches connected via E\_Ports to the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade. When you install a DC Dir Switch MP Extension Blade in a DC SAN Backbone Director or DC04 SAN Director, all devices connected to the SAN Director Fibre Channel ports are part of the backbone fabric.
- **Edge fabric**—A fabric that is attached to an EX\_Port on a 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade.

### 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade hop count

The 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade are counted in the same way as a Fibre Channel switch when determining the fabric hop count ([Figure 51 \(page 111\)](#)). Devices communicating across fabrics through an MP Router must adhere to both the B-series 7-hop limit within an edge fabric and the MP Router 12-hop limit.

### 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade usage

Use a DC Dir Switch MP Extension Blade:

- If you want to add multiple fabrics to a Meta SAN, and one of those fabrics contains a DC SAN Backbone Director or DC04 SAN Director with an available slot for a DC Dir Switch MP Extension Blade

---

**NOTE:** All devices connected to one of the SAN Director's Fibre Channel ports are part of the backbone fabric.

---

- If you want to share up to 256 local devices in a backbone fabric connected to multiple edge fabrics
- For FCIP, if there are DC SAN Backbone Directors or DC04 SAN Directors at both the local site and remote site

Use a 1606 Extension SAN Switch:

- If you do not have a DC SAN Backbone Director or DC04 SAN Director, and do not plan to use one
- For FCIP, if there are no DC SAN Backbone Directors or DC04 SAN Directors at the local site or remote site
- If you do not have an available slot in your DC SAN Backbone Director or DC04 SAN Director to install a DC Dir Switch MP Extension Blade

For use-case configuration examples, see the 400 MP Router and MP Router Blade use-case white papers at [http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg_R1002_USEN).

## Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules

This section describes the fabric rules for 8 Gb/s switches with license-enabled integrated Fibre Channel routing, the 400 MP Router, and the MP Router Blade, and other factors you should consider when building B-series fabrics that contain Fibre Channel routing. The fabric rules for Fibre Channel routing apply to SANs that include the same operating system and storage products as the B-series switches. See [“Operating systems and storage products” \(page 95\)](#). For information about configuring MP Routers for FCIP, see [“B-series 400 MP Router and MP Router Blade” \(page 283\)](#).

[Table 33 \(page 103\)](#) describes the rules for creating fabrics with Fibre Channel routing using 8 Gb/s switches, 400 MP Routers, and MP Router Blades.

**Table 33 Integrated Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules**

Rule number	Description
1	The MP Router Blade supports only the DC SAN Backbone Director, DC04 SAN Director, and SAN Director 4/256. The maximum number of MP Router Blades per DC SAN Backbone Director is four; the maximum number of MP Router Blades per DC04 SAN Director is four; the maximum number of MP Router Blades per SAN Director 4/256 is two.
2	All configurations must use the default setting for R_A_TOV (10,000 milliseconds) and E_D_TOV (2,000 milliseconds).
3	Devices connected to edge fabrics can be routed through an LSAN to a backbone fabric. <b>Note:</b> Edge fabric-to-backbone fabric routing is supported using the 400 MP Router and MP Router Blade only. It is not supported if an MP Router is configured in the backbone.
4	Direct connection of devices to the 400 MP Router or MP Router Blade Fibre Channel ports (F_Port, FL_Port) is supported.
5	Fibre Channel server boot through the 400 MP Router or MP Router Blade is not supported.
6	Each 400 MP Router or MP Router Blade can only attach to one backbone fabric; multiple routers can attach to the same backbone fabric.
7	If an MP Router Blade has any port configured as an EX_Port, then the blade Fibre Channel ports, all Fibre Channel ports in the director the blade is in, and all other switches connected to the director using an E_Port are part of the backbone fabric.
8	A backbone fabric supports a maximum of 256 devices.
9	For the 400 MP Router and MP Router Blade, you can configure a Fibre Channel port as an F_Port, FL_Port, E_Port, or EX_Port (Fibre Channel routing), and configure a GbE port as a VE_Port (FCIP) or VEX_Port (FCIP with Fibre Channel routing). Traffic switching is supported between the following port types: <ul style="list-style-type: none"> <li>• Fibre Channel switching (F_Port, FL_Port, E_Port)</li> <li>• Fibre Channel switching and Fibre Channel routing (F_Port, FL_Port, EX_Port)</li> <li>• Fibre Channel switching and FCIP (F_Port, FL_Port, VE_Port, VEX_Port)</li> <li>• Fibre Channel routing and FCIP (VE_Port, VEX_Port)</li> </ul> <b>Note:</b> VEX_Ports and EX_Ports cannot be connected to the same edge fabric.
10	FCIP is supported between pairs of 400 MP Routers, MP Router Blades, or between a 400 MP Router and MP Router Blade. FCIP is not supported between an MP Router and either a 400 MP Router or MP Router Blade. See <a href="#">“B-series 400 MP Router and MP Router Blade” (page 283)</a> .
11	FCIP is supported with IPsec data encryption. This requires a minimum firmware version of 5.2.0a. IPsec is supported with HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. <b>Note:</b> IPsec and FCIP FastWrite are mutually exclusive and cannot be configured simultaneously.

**Table 33 Integrated Fibre Channel routing, 400 MP Router, and MP Router Blade fabric rules** (continued)

Rule number	Description
12	<p>FCIP FastWrite support:</p> <ul style="list-style-type: none"> <li>FCIP is supported with FCIP FastWrite SCSI write acceleration. This requires a minimum firmware version of 5.2.0a. For firmware version 6.3.x, version 6.3.0c (or later) is required for FCIP FastWrite. FCIP FastWrite is supported with HP P9000 (XP) Continuous Access and is not supported with HP P6000 Continuous Access.</li> </ul> <p><b>Note:</b> FCIP FastWrite and IPsec are mutually exclusive and cannot be configured simultaneously.</p>
13	<p>Fibre Channel FastWrite support:</p> <ul style="list-style-type: none"> <li>FC FastWrite is supported starting with Fabric OS 5.3.0. However, the following rules apply while configuring FCFW:</li> <li>FCFW disables the local Ethernet ports (ge0 and ge1), making it impossible to configure FCFW and FCIP tunnels on the same 400 MPR or the MP Router blade.</li> <li>FCFW does not work in FICON environments.</li> <li>FCFW flows may be routed to another 400 MPR or MP Router blade on the FC network. This 400 MPR or MP Router blade may have active FCIP tunnels over an IP network. FCFW flows may be passed through the FCIP tunnel, but only if the FCIP FastWrite option is disabled on the tunnel.</li> <li>FCFW does not support loop device configurations for more than one device.</li> </ul> <p>For more information, see the <i>Fabric OS FCIP Administrator's Guide</i>, V6.3.0 (or later).</p>
14	<p>For redundancy, HP recommends using a minimum of two EX_Ports (router IFLs) for each fabric connected to an 400 MP Router or MP Router Blade. The number of EX_Ports required is based on the performance requirements for all devices shared by fabrics through the router.</p> <ul style="list-style-type: none"> <li>For initial settings, calculate the number of EX_Ports using a 5:1 ratio of standard ISLs or switch device ports to router EX_Port IFLs. Monitor port performance to determine if this ratio is acceptable based on usage. Use the <code>portperfshow</code> command to monitor router port performance.</li> <li>For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1.</li> </ul>
15	<p>Supports routing between B-series secured and nonsecured edge fabrics using Secure Fabric OS.</p>
16	<p>Supports routing between B-series and M-series/McDATA edge fabrics. See <i>HP StorageWorks Fabric Interoperability: Merging EX_Ports Based on M-series and B-series Fibre Channel Switches Application Notes</i> at <a href="http://hp.com/go/sdgmanuals">http://hp.com/go/sdgmanuals</a>.</p> <p><b>Note:</b> Continuous Access products are not supported in routed fabrics using both B-series and M-series/McDATA edge fabrics.</p>

## Scalability rules

Table 34 (page 105) lists the scalability rules for Meta SANs using integrated Fibre Channel routing, the 400 MP Router, and the MP Router Blade. For MP Router scalability rules, see “Scalability rules” (page 110). For configurations with an MP Router and a 400 MP Router or MP Router Blade, apply the rules in Table 34 (page 105).

The following terms describe MP Router scalability:

- **Front phantom domains**—All EX\_Ports on a Fibre Channel router that are connected to the same edge fabric are presented as a single front phantom domain.
- **Translate phantom domains**—Assigned to each remote fabric connected to a 400 MP Router or MP Router Blade. Each fabric has one unique translate domain, regardless of how many EX\_Ports are connected to the fabric.
- **Backbone fabric**—The 400 MP Router or MP Router Blade backbone fabric consists of at least one 400 MP Router or MP Router Blade and includes all MP Routers and B-series switches connected via E\_Port to the 400 MP Router or MP Router Blades. When you install an MP



Router Blade in a 4/256 SAN Director, all devices connected to the SAN Director Fibre Channel ports are part of the backbone fabric.

- **Edge fabric**—A fabric that is attached to an EX\_Port on a 400 MP Router or MP Router Blade.

**Table 34 Integrated Fibre Channel routing, 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, and MP Router Blade scalability rules for firmware 5.x/6.x/7.x**

Edge fabric scalability		
Domains <sup>1</sup>	Maximum number of front domains per edge fabric	10
	Maximum number of translate domains (number of remote fabrics) per edge fabric	31 47 (7.x)
	Maximum number of real domains (local switches) per edge fabric	26
	Maximum number of domains per edge fabric (real domains + front domains + translate domains) <b>Note:</b> The total number of domains in the first three rows must not exceed 69.	67 83 (7.x)
Devices (user ports)	Maximum number of local and remote devices per edge fabric. (For edge fabrics with more than 600 devices, 4.2.0c (or later) is required for all switches in the fabric when using the 400 MP Router, MP Router Blade, or MP Router.)	1,500 2,000 (with reduced number of max edge fabrics)
	Maximum number of imported devices per edge fabric	1,000
LSAN zone scalability		
Zoning	Maximum number of entries per LSAN zone <b>Note:</b> LSAN zones support WWN-based zoning only.	64
Meta SAN scalability		
Edge fabrics	Maximum number of edge fabrics connected in a Meta SAN	Edge fabrics containing up to 1,500 WWNs: <ul style="list-style-type: none"> <li>• 48 (DC and SN8000B Directors)</li> <li>• 24 (all others)</li> </ul> Edge fabrics containing up to 2,000 WWNs: <ul style="list-style-type: none"> <li>• 16 (DC and SN8000B Directors)</li> <li>• 16 (all others)</li> </ul>
Meta SAN	Maximum number of total ports per Meta SAN	12,800
	Maximum number of edge fabrics per chassis: MP Router and MP Router Blade	12
	Maximum number of edge fabrics per chassis: DC Director, 8/80 and 8/40 switches with Integrated Routing HP SN8000B 8-Slot SAN Backbone Director Switch, HP SN8000B 4-Slot SAN Director Switch	24 <sup>2</sup>
	Maximum number of EX_Ports per chassis with Integrated Routing (DC and SN8000B Directors, 8/80, 8/40 switches)	128/80/40
Routers	Maximum number of MP Routers per Meta SAN	10
LSAN	Maximum number of LSAN device entries (proxy devices) per Meta SAN	10,000
	Maximum number of LSAN zones per Meta SAN	2,500 (5.x)

**Table 34 Integrated Fibre Channel routing, 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, and MP Router Blade scalability rules for firmware 5.x/6.x/7.x (continued)**

		5,000 (6.x) 5,000 (7.x)
<b>Backbone fabric scalability</b>		
Switches	Maximum number of local switches per backbone fabric	5 (5.x) 12 (6.x) 12 (7.x)
Domains	Maximum number of translate domains per backbone fabric	32 48 (7.x)
	Maximum number of total domains per backbone fabric	42 60 (7.x)
WWNs	Maximum number of local WWNs per backbone fabric	512 (32 edge fabrics) 1,280 (1 edge fabric)
<b>Hop count scalability</b>		
Hop count	Maximum number of hops between switches (including routers) in a Meta SAN (Figure 51 (page 111))	12

<sup>1</sup> For more information, see Figure 52 (page 111).

<sup>2</sup> Tested to 24, but can support up to 48.

#### 400 MP Router and MP Router Blade hop count

The 400 MP Router and MP Router Blade are counted in the same way as a Fibre Channel switch when determining the fabric hop count (“MP Router with seven hops” (page 111)). Devices communicating across fabrics through an MP Router must adhere to both the B-series 7-hop limit within an edge fabric and the MP Router 12-hop limit.

#### 400 MP Router and MP Router Blade usage

Use an MP Router Blade:

- If you want to add multiple fabrics to a Meta SAN, and one of those fabrics contains a 4/256 SAN Director, DC SAN Backbone Director, or DC04 SAN Director with an available slot for an MP Router Blade

---

**NOTE:** All devices connected to the SAN Director Fibre Channel ports are part of a backbone fabric.

---

- If you want to share up to 256 local devices in a backbone fabric with multiple edge fabrics
- For FCIP, if there are 4/256 SAN Directors, DC SAN Backbone Directors, or DC04 SAN Directors at both the local site and remote site

Use a 400 MP Router:

- If you do not have a 4/256 SAN Director, DC SAN Backbone Director, or DC04 SAN Director, and do not plan to use one
- If you have a 4/256 SAN Director, DC SAN Backbone Director, or DC04 SAN Director, and do not want to use it as a backbone fabric
- For FCIP, if there are no 4/256 SAN Directors, DC SAN Backbone Directors, or DC04 SAN Directors at the local site or remote site

For use-case configuration examples, see the 400 MP Router and MP Router Blade use-case white papers at [http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg_R1002_USEN).

## MP Router fabric rules

This section describes the fabric rules for the MP Router and other factors you should consider when building B-series fabrics that contain MP Routers. The fabric rules for the MP Router apply to SANs that include the same operating system and storage products as the B-series switches. See “Operating systems and storage products” (page 95).

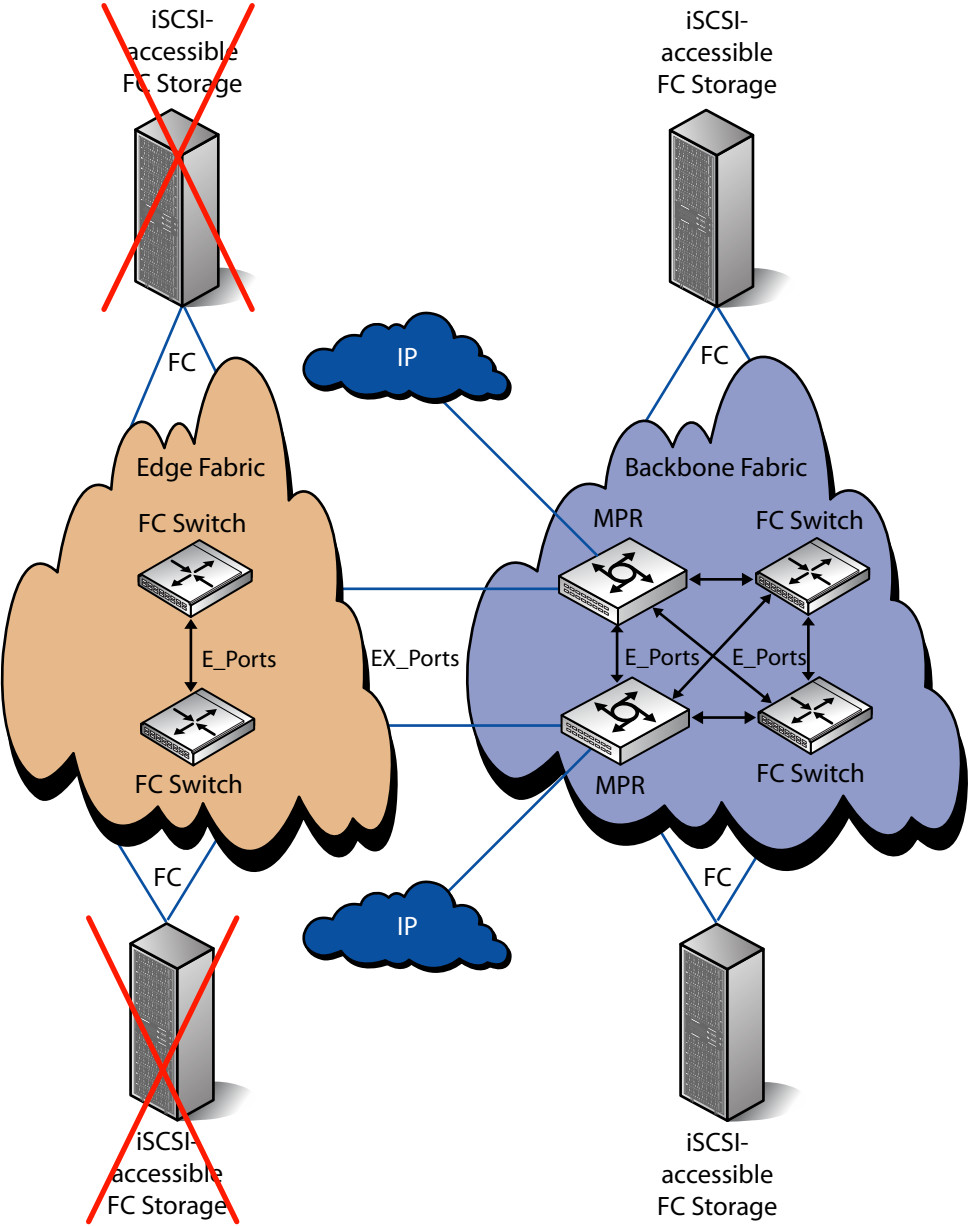
For information about configuring the MP Router for FCIP, see “B-series MP Router” (page 285).

Table 35 (page 107) describes the rules for creating fabrics with MP Routers.

**Table 35 MP Router fabric rules**

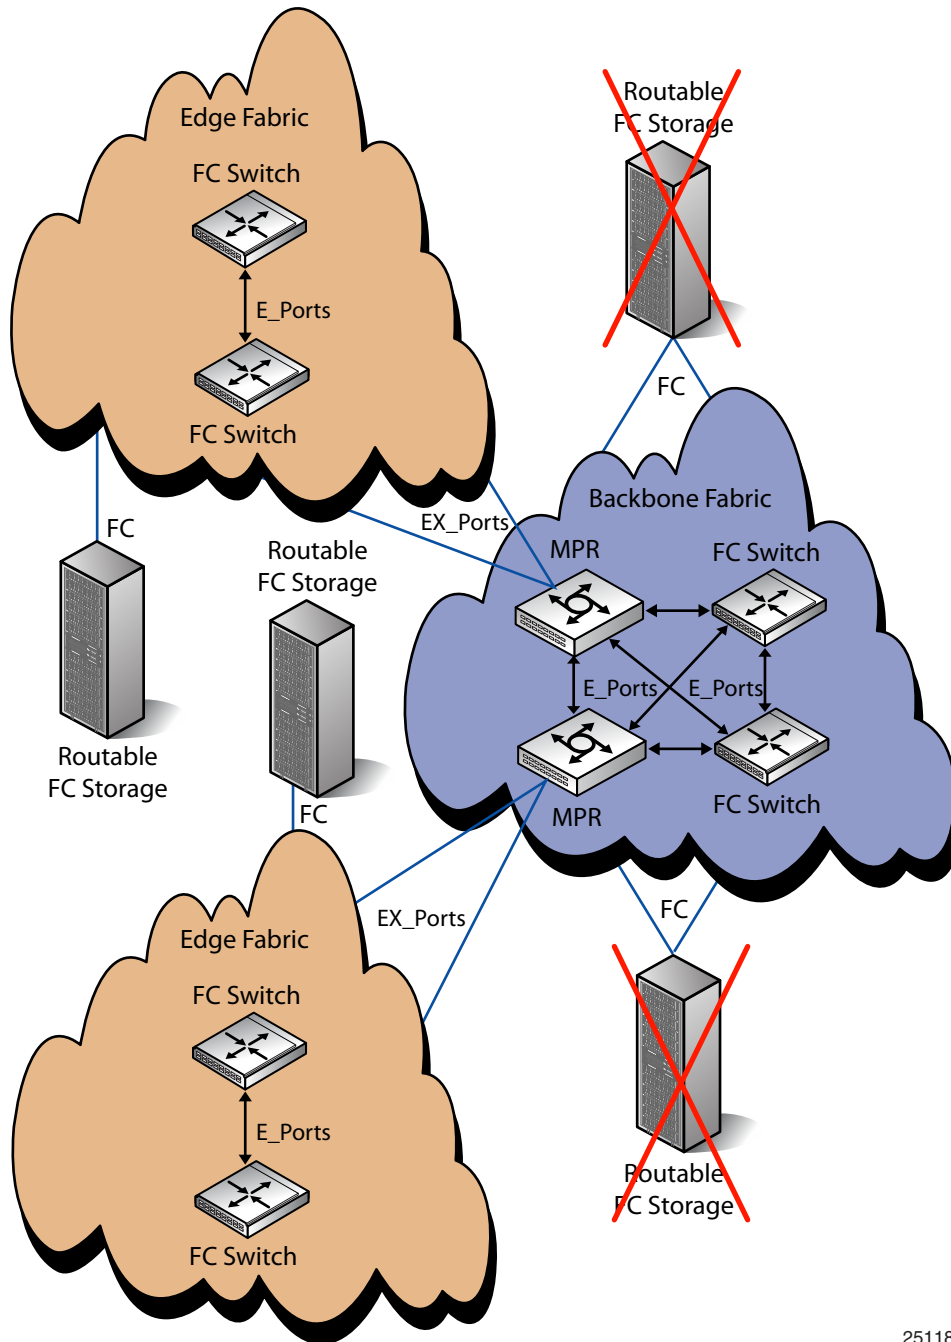
Rule number	Description
1	All configurations must use the default settings for R_A_TOV (10,000 milliseconds) and E_D_TOV (2,000 milliseconds).
2	Devices connected directly to the MP Router or the backbone fabric are for iSCSI access only and cannot be routed via an LSAN zone to an edge fabric (Figure 49 (page 108)).
3	Devices connected to edge fabrics are for local and routed access only and cannot be routed via an LSAN to an iSCSI port on the MP Router (Figure 50 (page 109)).
4	<p>The MP Router is not supported for use as a Fibre Channel switch. Direct device connectivity to MP Router ports (F_Port) is only supported with iSCSI. See “iSCSI storage” (page 312). You can configure a Fibre Channel port as an E_Port, EX_Port (Fibre Channel Routing), GbE iSCSI, or GbE FCIP. Traffic switching is supported between the following port types:</p> <ul style="list-style-type: none"> <li>• Fibre Channel and Fibre Channel routing (E_Port, EX_Port)</li> <li>• Fibre Channel and iSCSI (E_Port, F_Port, GbE iSCSI)</li> <li>• Fibre Channel and FCIP (E_Port, GbE FCIP)</li> <li>• Fibre Channel routing and FCIP (EX_Port, GbE FCIP)</li> </ul> <p>Traffic switching between other port combinations is not supported. This applies to port combinations on the same platform or independent platforms. For example, traffic from an iSCSI initiator connected to an MP Router cannot be routed to an EX_Port on the same router or independent router.</p>
5	Fibre Channel SAN boot or iSCSI server boot through the MP Router is not supported.
6	<p>For redundancy, HP recommends using a minimum of two EX_Ports (MP Router ISLs) for each fabric connected to the MP Router. The number of EX_Ports required is based on the performance requirements for all devices shared by fabrics through the MP Router.</p> <ul style="list-style-type: none"> <li>• For initial settings, calculate the number of EX_Ports using a 5:1 ratio of standard ISLs or switch device ports to MP Router EX_Port ISLs. Monitor port performance to determine if this ratio is acceptable based on usage. Use the <code>portperfshow</code> command to monitor MP Router port performance.</li> <li>• For very high bandwidth requirements and applications, use a ratio of 3:1 or 1:1.</li> </ul>
7	Supports routing between B-series secured and nonsecured edge fabrics using Secure Fabric OS. This requires MP Router firmware 7.4 (or later).
8	<p>Supports routing between B-series and M-series edge fabrics. This requires MP Router firmware 7.4 (or later), B-series firmware 2.6.2x, 3.2.x, or 5.x (or later), and M-series firmware 06.02.00-22 in Open Fabric mode. For MP Router firmware 7.4.1 (or later), M-series firmware 07.01.03 and 08.02.00 are supported in Open Fabric mode.</p> <p><b>Note:</b> Continuous Access products are not supported in routed fabrics using both B-series and M-series edge fabrics.</p>

Figure 49 iSCSI storage device connectivity for the MP Router



25117b

Figure 50 Fibre Channel routing storage device connectivity for the MP Router



25118a

The following sections describe additional considerations and rules for using the MP Router in a fabric:

- "XPath OS compatibility" (page 110)
- "Scalability rules" (page 110)
- "MP Router hop count" (page 111)
- "MP Router backbone fabric" (page 111)

## XPath OS compatibility

Table 21 (page 85) through Table 23 (page 87) list the B-series switches and firmware versions that are supported in a SAN with the MP Router XPath OS.

## Scalability rules

Table 36 (page 110) lists the scalability rules for Meta SANs using the MP Router.

The following terms describe MP Router scalability:

- **Front phantom domains**—Individual EX\_Port connections from the MP Router to the edge fabrics.
- **Translate phantom domains**—Assigned to each remote fabric connected to an MP Router. There is one unique translate domain for each fabric, regardless of how many EX\_Ports are connected to the fabric.
- **Backbone fabric**—The MP Router backbone fabric consists of at least one MP Router and includes all MP Routers and B-series switches connected to the MP Router via E\_Ports.
- **Edge fabric**—A fabric that is attached to an EX\_Port on an MP Router.

**Table 36 MP Router scalability rules**

Edge fabric scalability		
Domains <sup>1</sup>	Maximum number of front domains (number of EX_Ports) per edge fabric	24
	Maximum number of translate domains (number of remote fabrics) per edge fabric	33
	Maximum number of real domains (local switches) per edge fabric	34
	Maximum number of domains per edge fabric (real domains + front domains + translate domains) <b>Note:</b> The total number of domains in the first three rows must not exceed 80.	80
Devices (user ports)	Maximum number of local and remote devices per edge fabric. (For edge fabrics with more than 600 devices, 4.2.0c (or later) is required for all switches in the fabric when using the MP Router.)	1,280
	Maximum number of imported devices per edge fabric	1,000
LSAN zone scalability		
Zoning	Maximum number of entries per LSAN zone <b>Note:</b> LSAN zones support WWN-based zoning only.	200
Meta SAN scalability		
Edge fabrics	Maximum number of edge fabrics connected in a Meta SAN	34
Meta SAN	Maximum number of total ports per Meta SAN	12,800
Routers	Maximum number of MP Routers per Meta SAN	14
LSAN	Maximum number of LSAN device entries (proxy devices) per Meta SAN	10,000
	Maximum number of LSAN zones per Meta SAN	1,000
Hop count scalability		
Hop count	Maximum number of hops between switches (including routers) in a Meta SAN (Figure 51 (page 111))	12

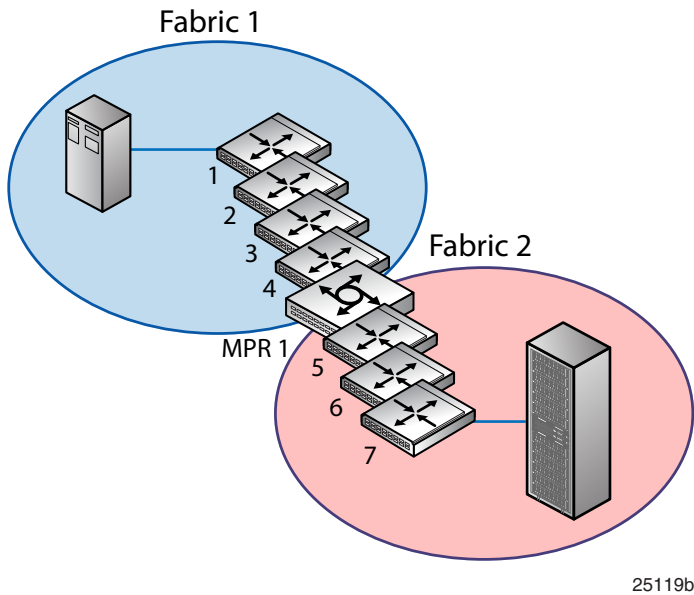
<sup>1</sup> For more information, see Figure 52 (page 111).

**NOTE:** For configurations that have an MP Router with a 400 MP Router or MP Router Blade in the same Meta SAN, apply the 400 MP Router and MP Router Blade scalability rules (lowest-common-denominator limits apply). See Table 34 (page 105).

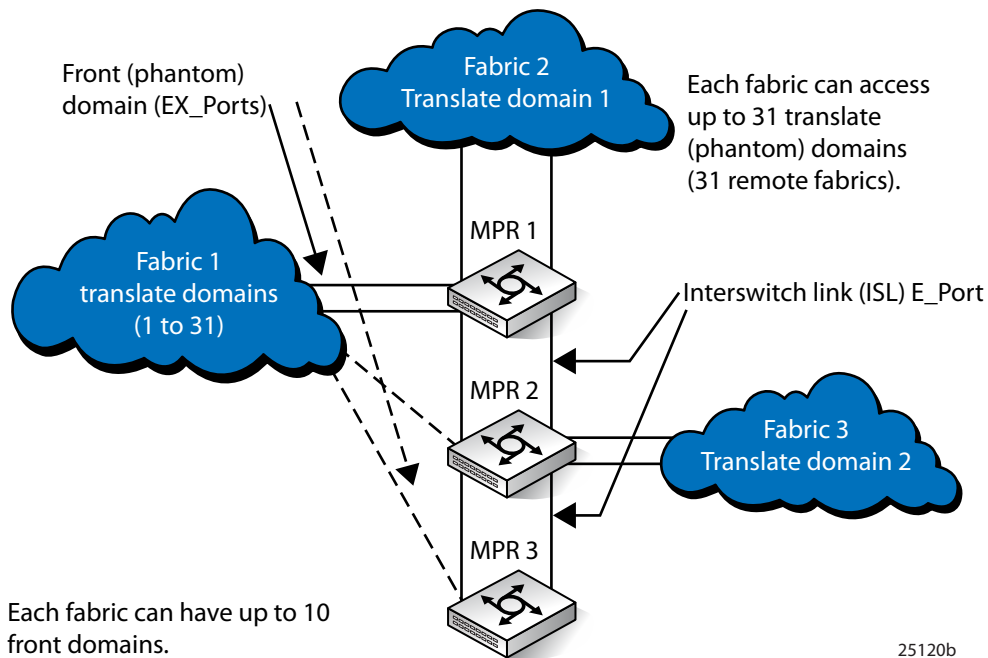
## MP Router hop count

The MP Router counts the same as a Fibre Channel switch is counted with respect to fabric hop count (Figure 51 (page 111)). Devices communicating across fabrics through an MP Router must adhere to both the B-series 7-hop limit within an edge fabric and the MP Router 12-hop limit.

**Figure 51 MP Router with seven hops**



**Figure 52 Front and translate domains**



## MP Router backbone fabric

An MP Router backbone fabric consists of at least one MP Router and includes all MP Routers and B-series Fibre Channel switches connected to the MP Router through E\_Ports. The MP Router backbone fabric is subject to all standard B-series Fibre Channel switch fabric rules. The MP Router counts as one Fibre Channel switch in terms of the maximum number of switches in a B-series fabric.

## Core switch addressing mode

B-series Fibre Channel switches using firmware 4.1 (or later) and all MP Router models are shipped with the Core switch PID parameter bit set to 1. Legacy switches using firmware 4.0 or 3.1 (or earlier) were shipped with this bit set to 0, which limited the number of switches in a fabric and the number of ports on a switch. HP recommends that you set the Core switch PID parameter bit to 1 in all B-series fabrics. Use the `configure` command to modify the Core switch PID setting.

---

**NOTE:** Beginning with Fabric OS 6.x, the only supported PID format is core pid (1). For more information, see the *HP StorageWorks Fabric OS Administrator Guide*.

---

If you change the Core switch addressing mode for a fabric, consider the following:

- If a B-series Fibre Channel fabric contains a Core Switch 2/64, SAN Director 2/128, SAN Director 4/256, or SAN Switch 4/8, 4/16, 4/32B, 4/32, or 4/64, you must set the Core switch PID parameter bit for all other switches in the fabric.
- All switches in the fabric must have the same Core switch PID parameter bit setting; otherwise, the fabric will segment.
- For multi-fabric SANs, you can change the Core switch PID parameter bit setting on one fabric at a time, allowing SAN operations to continue without interruption.
- The MP Router supports using different Core PID formats on different fabrics and routing frames between them.
- HP-UX and IBM AIX systems use the address bits to identify logical units. Therefore, if you change the Core switch PID setting, you must also change the logical unit definitions. After making these changes, you must reboot all servers in the SAN.
- If you add switches (other than the switches that require the Core switch PID settings) to a fabric in which the Core switch PID parameter bit is cleared, you must clear this bit on the new switches.

For more information about Core switch addressing mode, see the HP SAN Infrastructure website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

## Zoning limits and enforcement

Table 37 (page 113) describes zoning enforcement for B-series Fibre Channel switches and MP Router LSANs. See also “Zoning” (page 376).



**Table 37 Zoning enforcement for B-series Fibre Channel switches and MP Router LSANs**

Switches or routers	Configuration	Enforcement	Comments
Encryption SAN Switch SN8000B 8-Slot SAN Backbone Director SN8000B 4-Slot SAN Director SN6000B 16Gb FC Switch	Define zones using domain number and port number	Access authorization at frame level in hardware	Hard zoning
DC SAN Backbone Director DC04 SAN Director	Define zones using WWNs only		
SAN Director 2/128, 4/256 SAN Switch 8/8, 8/24, 8/40, 8/80 EVA4400 Embedded Switch Module, 8 Gb Brocade Brocade 8Gb SAN Switch for HP BladeSystem c-Class Brocade 4Gb SAN Switch for HP c-Class BladeSystem Brocade 4Gb SAN Switch for HP p-Class BladeSystem SAN Switch 4/8, 4/16, 4/32B, 4/32, 4/64 1606 Extension SAN Switch DC Dir Switch MP Extension Blade 400 MP Router, MP Router Blade, MP Router <sup>1</sup> B-series FCoE CN switch	Define zones using combination of domain/port numbers and WWNs	Name Server directory-based authentication, login authentication	Soft zoning, Name Servers discovery-based authentication, and login protection

<sup>1</sup> 400 MP Router, MP Router Blade, and MP Router LSAN zones support WWN-based zoning only.

## Zoning guidelines for B-series Fibre Channel switches

To configure B-series Fibre Channel switch zoning, observe the following best practices:

- For 4 Gb/s and 2 Gb/s SAN fabric switches, avoid transitions to soft zoning in a hardware-enforced zoning environment.

B-series Fibre Channel switches allow a maximum of 64 SID entries for each quad. If you exceed this limit, the affected ports transition from hard to soft enforcement. Although the switch logs display warning messages, data integrity is preserved during this transition.

- Maintain data access as defined in your SAN design, but avoid configuring hosts and targets on the same quad.
- Configure each quad with members of the same zone. Avoid configuring members of different zones on the same quad. For example, configure UNIX zone members on one quad and Windows members on a different quad.
- Minimize zone entries by including only hosts and targets that communicate. For example, rather than combine all hosts of the same OS type into one zone, make smaller zones with only hosts and targets that need to communicate.
- Use the `portzonestatus` command to display and verify the zoning status of each port.

The `portzonestatus` command displays the status of each port:

- **Hard**—Hardware enforcement
- **Soft**—Name server and ASIC-assisted authentication
- **All**—No zoning enforcement

### Zoning guidelines

The following messages indicate that a port has changed to soft zoning:

```
WARNING ZONE-ZONEGROUPADDFAIL, 3, WARNING - port 7 Out of CAM entries
```

WARNING ZONE-SOFTZONING, 3, WARNING - port 7: zoning enforcement changed to SOFT

The zoning configuration has exceeded limits, forcing the specified port to change from hardware-enforced zoning to software-enforced zoning. Other zone members remain hardware enforced. These warning messages appear at zoning configuration time (for port-level zoning) or dynamically at run time (for WWN zoning).

### Primary management switch recommendations (B-series Fibre Channel switches)

Recommendations for primary management switches follow:

- Designate one switch in the fabric as the primary management switch and use it for all management and control, including zoning, Time Services, Fabric Manager, and WebTools. Using one switch for access prevents multiple administrators from making changes to switches in the fabric at the same time. If you have a core-edge topology, HP recommends you use a core switch as the primary management switch. Typically, a core switch is connected directly to all other switches in a core-edge fabric, providing optimal communication. In SANs with mixed Fabric OS, these functions must only be executed from a switch running the highest Fabric OS. For example, if a SAN contains switches running Fabric OS 3.x and switches running Fabric OS 5.x, these functions must be executed from a switch running Fabric OS 5.x.
- Configure the primary management switch in the fabric as the preferred principal switch by using the `fabricprincipal` command.
- How a principal switch is assigned can vary in a fabric. That assignment depends on the state of the fabric, the switch WWN, and whether other switches or fabrics merge with that fabric. The principal switch in a fabric may not remain a principal switch once new switches are added to the original fabric or the original fabric is reconfigured. The Fibre Channel standards alone specify the mechanisms to implement the `fabricprincipal` command. These mechanisms allow a preference for a switch that requests to be the principal switch in a fabric. However, they do not absolutely guarantee that a switch that requests to be the principal switch actually achieves that status.

The primary management switch is used by Fabric Manager as the main access point to that fabric.

- Fabric time synchronization is implemented differently based on the security setting. For switches in a fabric where security is not enabled, synchronize time with the principal switch in the fabric. The principal switch in the fabric synchronizes its clock with an NTP time server by identifying the time server with the `tscklockserver` command. In a fabric where security is enabled, switches synchronize time with the primary Fibre Channel switch, which may or may not be the principal switch.

---

## 6 C-series switches and fabric rules

This chapter describes the C-series Fibre Channel and FCoE switches and the fabric rules for building C-series fabrics. It describes the following topics:

- “C-series Fibre Channel switches” (page 115)
- “Fibre Channel switch fabric rules” (page 126)
- “C-series FCoE Converged Network switches” (page 76)

---

**NOTE:** For information about using switches from the different series in the same SAN or fabric, see “SAN fabric connectivity and switch interoperability rules” (page 146).

---

### C-series Fibre Channel switches

C-series Fibre Channel switches:

- Can be core or edge switches. When configured in a core-edge topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- Are supported for use with FCIP through the HP IP Distance Gateway. For more information, see “SAN extension” (page 249).
- MDS 9506, MDS 9509, and MDS 9513 Director switches can accommodate various Fibre Channel Switching Modules, IP Storage Services Modules (IPS), a Multiprotocol Services Module, a Multiservice Module, and a Storage Services Module (SSM).
  - The 16-port 2 Gb/s Fibre Channel Switching Module is recommended for enterprise-level host connections, storage connections, and ISL connections.
  - The 32-port 2 Gb/s Fibre Channel Switching Module uses 3.2:1 internal oversubscription and is recommended for low- to mid-range host connections and tape device connections.
  - The 12-port 4 Gb/s Fibre Channel Switching Module is recommended for high-performance 4 Gb/s enterprise-level host connections, storage connections, and ISL connections.
  - The 24-port 4 Gb/s Fibre Channel Switching Module uses 2:1 internal oversubscription across four 6-port port groups and is recommended for 4 Gb/s enterprise-level mid-range host and storage connections.
  - The 48-port 4 Gb/s Fibre Channel Switching Module uses 4:1 internal oversubscription across four 12-port port groups and is recommended for 4 Gb/s low-range host connections and tape devices connections.
  - The 4-port 10 Gb/s Fibre Channel Switching Module is ideal for ultra-high-bandwidth ISL connectivity and inter-data center connection over a metropolitan optical infrastructure.
  - IPS modules, Multiprotocol Services Modules, and Multiservice Modules provide MDS iSCSI and FCIP functionality:
    - IPS-4 and IPS-8 provide 4 and 8 GbE IP ports, respectively.
    - The 14/2 Multiprotocol Services Module provides 14 2 Gb/s Fibre Channel ports and 2 GbE IP ports.
    - The 18/4 Multiservice Module provides 18 4 Gb/s Fibre Channel ports and 4 GbE IP ports.
  - The 32-port 2 Gb/s SSM is recommended for Fibre Channel write acceleration and SCSI statistics.

- The 48-port 8 Gb/s Host-Optimized Fibre Channel Switching Module uses four 12-port port groups at 12.8 Gb/s per port group and is recommended for 8 Gb/s low-range host connections.
- The 48-port 8 Gb/s Performance Fibre Channel Switching Module uses eight 6-port port groups at 12.8 Gb/s per port group and is recommended for 8 Gb/s enterprise-level mid-range host and storage connections.
- The 24-port 8 Gb/s Fibre Channel Switching Module uses eight 3-port port groups at 12.8 Gb/s per port group and is recommended for high-performance 8 Gb/s enterprise-level host connections, storage connections, and ISL connections.
- MDS 9216A switches have 2 slots:
  - One slot is a fixed configuration with a 16-port 2 Gb/s Fibre Channel Switching Module.
  - The second slot can accommodate a 16-port or 32-port 2 Gb/s Fibre Channel Switching Module; 12, 24, or 48-port 4 Gb/s Fibre Channel Switching Module; 4-port 10 Gb/s Fibre Channel module; or IPS-4, IPS-8, 14/2 Multiprotocol Services Module, or 18/4 Multiservice Module for iSCSI and FCIP support, and the 32-port 2 Gb/s SSM.
- MDS 9216i switches have 2 slots:
  - One slot is a fixed configuration with a 14/2 Multiprotocol Services Module.
  - The second slot can accommodate a 16-port or 32-port 2 Gb/s Fibre Channel Switching Module; 12, 24, or 48-port 4 Gb/s Fibre Channel Switching Module; 4-port 10 Gb/s Fibre Channel Switching Module; IPS-4, IPS-8, 14/2 Multiprotocol Services Module, or 18/4 Multiservice Module for iSCSI and FCIP support; and the 32-port 2 Gb/s SSM.
- MDS 9222i switches have 2 slots:
  - One slot is a fixed configuration with a 18/4 Multiservice Module.
  - The second slot can accommodate a 48-port 8 Gb/s Host-Optimized Fibre Channel Switching Module; 12, 24, or 48-port 4 Gb/s Fibre Channel Switching Module; 4-port 10 Gb/s Fibre Channel Switching Module; IPS-8 or 18/4 Multiservice Module for iSCSI and FCIP support; and the 32-port 2 Gb/s SSM.
- MDS 9134 switches have a base 34-port chassis to accommodate up to 32 4 Gb/s and 2 10 Gb/s ports. Four configurations are available:
  - 24 4 Gb/s active ports with optional 8-port 4 Gb/s and 2-port 10 Gb/s software expansion licenses
  - 32 4 Gb/s active ports with an optional 2-port 10 Gb/s software expansion license
  - 24 4 Gb/s and 2 10 Gb/s active ports  
A 48-port stackable solution is available with the purchase of two MDS 9134 switches.
  - 32 4 Gb/s and 2 10 Gb/s active ports  
A 64-port stackable solution is available with the purchase of two MDS 9134 switches.
- MDS 9124 switches have fixed configurations with 8, 16, and 24 full-rate 4 Gb/s Fibre Channel ports. The 8 and 16-port switches can be upgraded in 8-port increments with a software expansion license.
- MDS 9124e 12-port Fabric switch and the MDS 9124e 24-port Fabric switch for the c-Class BladeSystem have fixed configurations with 8 internal/4 external or 16 internal/8 external full-rate 4 Gb/s Fibre Channel ports. The 12-port switch can be upgraded to a 24-port switch with a software expansion license.
- MDS 8Gb 12-port fabric switch and the MDS 8Gb 24-port fabric switch for HP BladeSystem c-Class have fixed configurations with 8 internal/4 external or 16 internal/8 external full-rate

8 Gb/s Fibre Channel ports. The 12-port switch can be upgraded to a 24-port switch with a software expansion license.

- MDS 9020/9120/9140 switches have fixed configurations with 20 and 40 Fibre Channel ports, respectively:
  - MDS 9120 has 4 full-rate Fibre Channel ports and 16 2 Gb/s oversubscribed ports.
  - MDS 9140 has 8 full-rate Fibre Channel ports and 32 2 Gb/s oversubscribed ports.
  - MDS 9020 has a fixed configuration with 20 full-rate 4 Gb/s Fibre Channel ports.
- SN6000C Fabric Switches have fixed configurations with 16 and 32 full-rate 8 Gb/s Fibre Channel ports. The 16 and 32-port switches can be upgraded in 8-port increments with a software expansion license.

The C-series Fibre Channel switches offer:

- High availability
- Scalability
- Cost efficiency

## Model naming

The C-series Fibre Channel switches are named MDS 9 $xnn$ . The 95 $nn$  switches (Multilayer Directors) are director or core switches. The  $nn$  value indicates the number of slots available for supervisors and port modules. The 92 $nn$  switches are mid-range switches; the 91 $nn$  switches are entry-level switches; the MDS 9020 switch is an entry-level 4 Gb/s switch (Multilayer Fabric switches). For the 92 $nn$ , 91 $nn$ , and 90 $nn$ , the  $nn$  value indicates the number of fixed ports.

Multiprotocol and Multiservice products are designated with the  $i$  suffix, such as the MDS 9216 $i$  and MDS 9222 $i$ , or as IP storage services modules, such as the IPS-4, IPS-8, 14/2, and 18/4. For IPS products, the number indicates the total number of IP ports available. For example, the IPS-4 has 4 IP ports and the IPS-8 has 8 IP ports.

## Switch models

---

**NOTE:** Starting with release 4.1(1b), MDS 9000 SAN-OS software is rebranded to MDS 9000 NX-OS software. Software prior to NX-OS 4.1(1b) uses the SAN-OS name and remains on the 3.x code stream. NX-OS is designed to operate with switches running SAN-OS software.

---

Table 38 (page 118), Table 39 (page 118), and Table 40 (page 119) describe the C-series Fibre Channel switches, and Table 41 (page 119) describes the legacy C-series Fibre Channel switches.

HP supports all C-series Fibre Channel switches in a fabric if you:

- Use the firmware versions listed in Table 38 (page 118), Table 39 (page 118), Table 40 (page 119), and Table 41 (page 119).

When using NX-OS, SAN-OS, and FabricWare in the same fabric, HP strongly recommends upgrading all fabric switches in the SAN to the latest supported software version. All switches in the same major NX-OS/SAN-OS/FabricWare family must use the same software version. When updating switch firmware, you can use two successive NX-OS/SAN-OS/FabricWare versions temporarily. Do not enable new switch features until the upgrade is complete.

- Follow the fabric rules. See “Fibre Channel switch fabric rules” (page 126).

For the latest information on supported C-series Fibre Channel switches and firmware versions, see the HP SAN Infrastructure website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

**Table 38 C-series Fibre Channel switches for SAN-OS**

Switch	Firmware version	Maximum number of Fibre Channel ports
MDS 9506 Multilayer Director	SAN-OS 3.3(5a), 3.3(5), 3.3(4a), 3.3(4), 3.3(3), 3.3(2), 3.3(1c), 3.3(1a), 3.2(3), 3.2(2c), 3.2(1a), 3.1(3a), 3.1(3), 3.1(2a)	192
MDS 9509 Multilayer Director		336
MDS 9513 Multilayer Director		528
MDS 9222i Multiservice Fabric		66
MDS 9124e 12-Port Fabric Switch for HP c-Class BladeSystem		8 internal, 4 external
MDS 9124e 24-Port Fabric Switch for HP c-Class BladeSystem		16 internal, 8 external

**NOTE:** MDS 9509 requires Supervisor-2 or Supervisor-2A Modules to use the maximum port count.

**Table 39 C-series Fibre Channel switches for NX-OS 4.x**

Switch	Firmware version	Maximum number of Fibre Channel ports
MDS 9506 Multilayer Director	NX-OS 4.2(7d), 4.2(7b), 4.2(7a), 4.2(5), 4.2(3), 4.2(1a), 4.1(3a), 4.1(1c), 4.1(1b)	192
MDS 9509 Multilayer Director		336
MDS 9513 Multilayer Director		528
MDS 9222i Multiservice Fabric		66
MDS 9124e 12-Port Fabric Switch for HP c-Class BladeSystem		8 internal, 4 external
MDS 9124e 24-Port Fabric Switch for HP c-Class BladeSystem		16 internal, 8 external

**Table 40 C-series Fibre Channel switches for NX-OS 5.x**

Switch	Firmware version	Maximum number of Fibre Channel ports
SN6000C Fabric Switch	NX-OS 5.0(4b), 5.0(4), 5.0(1a)	48
MDS 9506 Multilayer Director		192
MDS 9509 Multilayer Director		336
MDS 9513 Multilayer Director		528
MDS 9222i Multiservice Fabric		66
MDS 9124e 12-Port Fabric Switch for HP c-Class BladeSystem		8 internal, 4 external
MDS 9124e 12-Port Fabric Switch for HP c-Class BladeSystem		16 internal, 8 external
MDS 8Gb 12-Port Fabric Switch for HP BladeSystem c-Class		8 internal, 4 external
MDS 8Gb 24-Port Fabric Switch for HP BladeSystem c-Class		16 internal, 8 external

**NOTE:** NX-OS requires Supervisor-2 or Supervisor-2A Modules for the MDS 950x Director switches.

**Table 41 C-series Fibre Channel legacy switches**

Switch	Firmware version	Maximum number of Fibre Channel ports
MDS 9020 Fabric	FabricWare 2.1(3), 2.1(2)	20
MDS 9120 Fabric	SAN-OS 3.3(5a), 3.3(5), 3.3(4a), 3.3(4), 3.3(3), 3.3(2), 3.3(1c), 3.3(1a), 3.2(3), 3.2(2c), 3.2(1a), 3.1(3a), 3.1(3), 3.1(2a)	20
MDS 9140 Fabric		40
MDS 9216A Multilayer Fabric		64
MDS 9216i Multilayer Fabric	NX-OS 4.2(7d), 4.2(7b), 4.2(7a), 4.2(5), 4.2(3), 4.2(1a), 4.1(3a), 4.1(1c), 4.1(1b) SAN-OS 3.3(5), 3.3(4a), 3.3(4), 3.3(3), 3.3(2), 3.3(1c), 3.3(1a), 3.2(3), 3.2(2c), 3.2(1a), 3.1(3a), 3.1(3), 3.1(2a)	62
MDS 9124 Fabric	NX-OS 5.0(4b), 5.0(4), 5.0(1a)	24
MDS 9134 Fabric		32 at 4 Gb/s 2 at 10 Gb/s

Table 42 (page 120) describes SAN-OS C-series switching module support.

**Table 42 C-series Fibre Channel switching module support matrix for SAN-OS**

Switching module	Switch type					
	MDS 9216A	MDS 9216i	MDS 9222i	MDS 9506 HP SN8000C 6-Slot SUP2A Director Switch	MDS 9509 HP SN8000C 9-Slot SUP2A Director Switch	MDS 9513 HP SN8000C 13-Slot SUP2A FAB2 Director Switch
16-port Fibre Channel Switching Module (1 Gb, 2 Gb)	Yes	Yes	No	Yes	Yes	Yes
32-port Fibre Channel Switching Module (1 Gb, 2 Gb)	Yes	Yes	No	Yes	Yes	Yes
14/2 Multiprotocol Services Module (1 Gb, 2 Gb, GbE)	Yes	Yes	No	Yes	Yes	Yes
18/4 Multiservice Module (1 Gb, 2 Gb, 4 Gb, GbE)	Yes	Yes	Yes	Yes	Yes	Yes
32-port Storage Services Module (1 Gb, 2 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
4-port IP Storage Services Module (GbE)	Yes	Yes	No	Yes	Yes	Yes
8-port IP Storage Services Module (GbE)	Yes	Yes	Yes	Yes	Yes	Yes
12-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
24-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
48-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
4-port Fibre Channel Switching Module (10 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
Supervisor-1 Switching Module	No	No	No	Yes	Yes	No
Supervisor-2 or Supervisor-2A Switching Module	No	No	No	Yes	Yes	Yes

**NOTE:** The MDS 9020, 9120, 9140, 9124, 9134, and 9124e switches have fixed configurations and do not support additional switching modules.

The 8 Gb/s Fibre Channel Switching Modules require NX-OS.



Table 43 (page 121) and Table 44 (page 122) describe NX-OS C-series switching module support.

**Table 43 C-series Fibre Channel switching module support matrix for NX-OS 4.x**

Switching module	Switch type					
	MDS 9216i	MDS 9222i	MDS 9506 with SUP2 or SUP2A Modules HP SN8000C 6-Slot SUP2A Director Switch	MDS 9509 with SUP2 or SUP2A Modules HP SN8000C 9-Slot SUP2A Director Switch	MDS 9513 with Fabric-1 Modules	MDS 9513 with Fabric-2 Modules HP SN8000C 13-Slot SUP2A FAB2 Dir Switch
16-port Fibre Channel Switching Module (1 Gb, 2 Gb)	Yes	No	Yes	Yes	Yes	Yes
32-port Fibre Channel Switching Module (1 Gb, 2 Gb)	Yes	No	Yes	Yes	Yes	Yes
14/2 Multiprotocol Services Module (1 Gb, 2 Gb, GbE)	Yes	No	Yes	Yes	Yes	Yes
18/4 Multiservice Module (1 Gb, 2 Gb, 4 Gb, GbE)	Yes	Yes	Yes	Yes	Yes	Yes
32-port Storage Services Module (1 Gb, 2 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
12-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
24-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
48-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
4-port Fibre Channel Switching Module (10 Gb)	Yes	Yes	Yes	Yes	Yes	Yes
4-port and 8-port IP Storage Services Module (GbE)	No	No	No	No	No	No
48-port 8Gb Host-Optimized Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	No	Yes	Yes	Yes	Yes	Yes
48-port 8Gb Performance Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	No	No	Yes	Yes	No	Yes
24-port 8Gb Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	No	No	Yes	Yes	No	Yes

**Table 44 C-series Fibre Channel switching module support matrix for NX-OS 5.x**

<b>Switching module</b>	<b>MDS 9222i</b>	<b>MDS 9506 with SUP2 or SUP2A Modules HP SN8000C 6-Slot SUP2A Director Switch</b>	<b>MDS 9509 with SUP2 or SUP2A Modules HP SN8000C 9-Slot SUP2A Director Switch</b>	<b>MDS 9513 with Fabric-1 Modules</b>	<b>MDS 9513 with Fabric-2 Modules HP SN8000C 13-Slot SUP2A FAB2 Dir Switch</b>
16-port Fibre Channel Switching Module (1 Gb, 2 Gb)	No	No	No	No	No
32-port Fibre Channel Switching Module (1 Gb, 2 Gb)	No	No	No	No	No
14/2 Multiprotocol Services Module (1 Gb, 2 Gb, GbE)	No	No	No	No	No
18/4 Multiservice Module (1 Gb, 2 Gb, 4 Gb, GbE)	Yes	Yes	Yes	Yes	Yes
32-port Storage Services Module (1 Gb, 2 Gb)	No	No	No	No	No
12-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes
24-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes
48-port Fibre Channel Switching Module (1 Gb, 2 Gb, 4 Gb)	Yes	Yes	Yes	Yes	Yes
4-port Fibre Channel Switching Module (10 Gb)	Yes	Yes	Yes	Yes	Yes
4-port and 8-port IP Storage Services Module (GbE)	No	No	No	No	No
48-port 8Gb Host-Optimized Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	Yes	Yes	Yes	Yes	Yes

**Table 44 C-series Fibre Channel switching module support matrix for NX-OS 5.x (continued)**

Switching module	MDS 9222i	MDS 9506 with SUP2 or SUP2A Modules HP SN8000C 6-Slot SUP2A Director Switch	MDS 9509 with SUP2 or SUP2A Modules HP SN8000C 9-Slot SUP2A Director Switch	MDS 9513 with Fabric-1 Modules	MDS 9513 with Fabric-2 Modules HP SN8000C 13-Slot SUP2A FAB2 Dir Switch
48-port 8Gb Performance Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	No	Yes	Yes	No	Yes
24-port 8Gb Fibre Channel Switching Module (2 Gb, 4 Gb, 8 Gb)	No	Yes	Yes	No	Yes

**NOTE:** The MDS 9020, 9120, 9140, 9124, 9134, 9124e, 8Gb fabric switch for HP BladeSystem c-Class, and SN6000C switches have fixed configurations and do not support additional switching modules.

The MDS 9513 requires Fabric-2 Modules to support the 48-port 8 Gb/s Performance Fibre Channel Switching Module and the 24-port 8 Gb/s Fibre Channel Switching Module. Fabric modules are also known as crossbar modules.

## Features

Features of the C-series Fibre Channel switches include:

- Nonblocking architecture using VOQ
- VSAN deployment over the physical infrastructure  
VSANs are separate instances of all fabric services, including address space.
- Advanced diagnostics and troubleshooting (FC Ping, FC Traceroute, SPAN, RSPAN, and Call Home)
- Comprehensive security (SSH, SFTP, RADIUS, SNMPv3, and RBAC)
- Comprehensive fabric management (CLI, SNMP, and Java-based GUI)
- Traffic management (FCC and QoS)
- High-availability, fault-tolerant software
- PortChannel (ISL aggregation for highly resilient SAN architectures)
- Integrated Multiprotocol capability (MDS 95nn and 92nn) for SAN extension (FCIP and iSCSI)
- NPV—The following C-series Fibre Channel switches are NPV compliant: MDS 9124, MDS 9134, MDS 9124e, 8Gb fabric switch for HP BladeSystem c-Class, and SN6000C. NPV requires SAN-OS 3.2(1a) (or later). See [“SAN fabric connectivity and switch interoperability rules” \(page 146\)](#).
- FlexAttach virtual pWWN

The Cisco MDS 9124e Fabric Switch for HP c-Class BladeSystem, MDS 9124, MDS 9134, 8Gb fabric switch for HP BladeSystem c-Class, and SN6000C support NPV with FlexAttach. FlexAttach provides automatic mapping of physical WWNs to virtual WWNs using NAT. When NPV mode is enabled, FlexAttach allows SAN and server administrators to install and replace servers without having to rezone or reconfigure the SAN. For more information on FlexAttach, see [“NPV with FlexAttach” \(page 168\)](#).

- Storage Media Encryption (SME)  
The MDS 9222i Multiservice Fabric Switch and the 18/4 Multiservice Module are SME compliant.
- IO Accelerator (IOA)  
The MDS 9222i Multiservice Fabric Switch and the 18/4 Multiservice Module are IOA compliant.

Additional features of the Director switches include:

- Nondisruptive software upgrades
- Hot-swappable line cards, supervisors, power supplies, and SFPs
- Redundant supervisor, cross-bar fabric, and power supplies

**NOTE:** The following features are not supported by MDS 9124, MDS 9124e, or MDS 9134:

- IVR
- Remote SPAN
- Translative loop support
- FCC—no generation, quench reaction only

Features of the C-series MDS 9020 running FabricWare include:

- Full nonblocking 4 Gb/s Fibre Channel performance on 20 line-rate ports
- Nondisruptive software upgrades
- Comprehensive fabric management (CLI, SNMP, and Java-based GUI)
- Compact, extremely cost-effective design

Table 45 (page 124) provides a comparison of the high-availability features for C-series Fibre Channel switches.

**Table 45 C-series Fibre Channel switch high-availability feature comparison**

Model	Redundant/ hot-swappable power	Redundant/ hot-swappable cooling	Redundant control processor	Nondisruptive code activation	Port module support	Protocol support
SN6000C Fabric	Yes/Yes	Yes/Yes	No	Yes	No	FC
MDS 9120/9140 Fabric	Yes/Yes	Yes/Yes	No	No	No	FC
MDS 9124/9134 Fabric	Yes/Yes	Yes/Yes	No	Yes	No	FC
MDS 9124e Fabric	N/A	N/A	No	Yes	No	FC
MDS 9216A/9216i Multilayer Fabric	Yes/Yes	Yes/Yes	No	No	Yes	FC FCIP iSCSI
MDS 9222i Multiservice Fabric	Yes/Yes	Yes/Yes	No	Yes	Yes	FC FCIP iSCSI
MDS 9506/9509/9513 Director	Yes/Yes	Yes/Yes	Yes	Yes	Yes	FC FCIP iSCSI

**Table 45 C-series Fibre Channel switch high-availability feature comparison** *(continued)*

Model	Redundant/hot-swappable power	Redundant/hot-swappable cooling	Redundant control processor	Nondisruptive code activation	Port module support	Protocol support
MDS 9020	No/No	Yes/No	No	Yes	No	FC
MDS 8Gb fabric switch for HP BladeSystem c-Class	N/A	N/A	No	Yes	No	FC

## Usage

Table 46 (page 125) describes the use of C-series Fibre Channel switches as core switches.

**Table 46 Using C-series Fibre Channel switches as core switches**

Model	1–48 total ports	49–224 total ports	225–512 total ports
SN6000C Fabric	Excellent (48-port stackable solution)	Not recommended	
MDS 9120 Fabric	Excellent (20 port maximum)		
MDS 9140 Fabric	Excellent (40 port maximum)		
MDS 9124 Fabric	Excellent (24 port maximum)		
MDS 9134 Fabric	Excellent (48-port stackable solution)	Excellent (64-port stackable solution)	Not recommended
MDS 9124e Fabric	Excellent (8 port maximum)	Not recommended	
MDS 9216A/9216i Multilayer Fabric MDS 9222i Multiservice Fabric	Excellent		
MDS 9506 Multilayer Director	Good	Excellent	
MDS 9509 Multilayer Director			
MDS 9513 Multilayer Director			
MDS 9020 Fabric	Excellent (20 port maximum)	Not recommended	
MDS 8Gb fabric switch for HP BladeSystem c-Class	Excellent (8 port maximum)		

Table 47 (page 126) describes the use of C-series Fibre Channel switches as edge switches.

**Table 47 Using C-series Fibre Channel switches as edge switches**

Model	1–48 total ports	49–224 total ports	225–512 total ports
SN6000C Fabric	Excellent		
MDS 9120 Fabric			
MDS 9140 Fabric			
MDS 9124 Fabric			
MDS 9134 Fabric			
MDS 9124e Fabric			
MDS 9216A/9216i Multilayer Fabric MDS 9222i Multiservice Fabric	Excellent	Very good	
MDS 9506 Multilayer Director	Excellent (select modules to optimize performance or user port count)		
MDS 9509 Multilayer Director			
MDS 9513 Multilayer Director			
MDS 9020	Excellent		
MDS 8Gb fabric switch for HP BladeSystem c-Class			

## Fibre Channel switch fabric rules

This section describes the fabric rules for C-series Fibre Channel switches and other factors you should consider when building C-series fabrics.

## Operating systems and storage products

The fabric rules for C-series switches apply to SANs that include the operating systems and storage products listed in [Table 48 \(page 126\)](#).

**Table 48 C-series Fibre Channel switch operating system and storage system support**

Operating systems	Storage products
<ul style="list-style-type: none"> <li>• HP-UX</li> <li>• OpenVMS</li> <li>• Tru64 UNIX</li> <li>• Apple Mac OS X</li> <li>• Citrix Xen</li> <li>• IBM AIX</li> <li>• Linux</li> <li>• Microsoft Windows</li> </ul>	<ul style="list-style-type: none"> <li>• P6300/P6500 EVA</li> <li>• EVA8400/6400</li> <li>• EVA4400</li> <li>• EVA8100/8000/6100/6000/4100/4000</li> <li>• EVA3000/5000</li> <li>• MSA2000fc G2 (MSA2300fc)/MSA2000fc</li> <li>• MSA1500/1000</li> <li>• P9500</li> </ul>

**Table 48 C-series Fibre Channel switch operating system and storage system support** (continued)

Operating systems	Storage products
<ul style="list-style-type: none"> <li>• Sun Solaris</li> <li>• VMware ESX</li> </ul>	<ul style="list-style-type: none"> <li>• P2000 G3 FC</li> <li>• VA7410/7400/7110/7100</li> <li>• XP24000/20000</li> <li>• XP12000/10000</li> <li>• XP1024/128</li> <li>• XP512/48</li> <li>• 3PAR F-Class, T-Class</li> <li>• HP P9000 (XP) Continuous Access</li> <li>• HP P6000 Continuous Access</li> </ul> <p>See “HP P6000 Continuous Access SAN integration” (page 220)</p>

The operating systems listed in Table 48 (page 126) may not be supported with every storage system listed. For current operating system and storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

See “Heterogeneous server rules” (page 159) and the following storage system rules chapters:

- “P2000 and MSA storage system rules” (page 203)
- “P6000/EVA storage system rules” (page 213)
- “P9000/XP and VA storage system rules” (page 227)

## Fabric rules for C-series Fibre Channel switches

The following fabric rules apply to all C-series Fibre Channel SANs. They also apply, in general, to HP P9000 (XP) Continuous Access and HP P6000 Continuous Access configurations. However, additional rules apply to HP Continuous Access implementations. For more information, see the *HP P6000 Continuous Access Implementation Guide*.

Table 49 (page 127) lists the rules for creating a SAN with C-series Fibre Channel switches.

**Table 49 C-series Fibre Channel switch fabric rules**

Rule number	Description
1	Up to 60 MDS switches with up to 4,000 total ports and 3,500 user ports in a fabric
2	MDS 9506 switch supports up to 192 ports over four modular slots (four 48-port modules <sup>1</sup> ).
3	MDS 9509 switch supports up to 336 ports over seven modular slots (seven 48-port modules <sup>1</sup> ). <b>Note:</b> This requires Supervisor-2 or Supervisor-2A Modules.
4	MDS 9513 switch supports up to 528 ports over eleven modular slots (eleven 48-port modules <sup>1</sup> ).
5	Maximum of seven switch hops (eight switches) between any two communicating devices
6	Maximum of 80 VSANs per fabric <sup>2</sup>
7	Maximum of 4,000 IVR shared devices per routed fabric <sup>3</sup> <b>Note:</b> VSANs and IVR are not supported on the MDS 9020. IVR is not supported on the MDS 9124, MDS 9124e, and MDS 9134.
8	Maximum of 2,000 IVR zones per routed fabric <sup>3</sup>
9	Persistent FCID mode (default setting) is required on all C-series Fibre Channel switches in fabrics or VSANs that include HP-UX or IBM AIX servers when implementing port-based fabric zoning.

**Table 49 C-series Fibre Channel switch fabric rules** (continued)

Rule number	Description
10	NPV mode with Cisco MDS 9124e for HP c-Class BladeSystem, MDS 9124, and MDS 9134 require SAN-OS 3.2(1a) (or later). NPV mode supports interoperability with B-series firmware 5.3.0, 5.3.0b, and 6.1.0d, and M-series firmware 09.06.02. See <a href="#">“BladeSystem with Cisco N_Port Virtualization mode”</a> (page 165).
11	FCIP is supported with IPsec data encryption using the 18/4 Multiservice Module, 14/2 Multiprotocol Services Module, MDS 9216i, and MDS 9222i. IPsec is supported with HP P6000 Continuous Access and HP P9000 (XP) Continuous Access.
12	Maximum of 114 FLOGIs or F Disc per port group for the 9200 and 9500 series switches. Maximum of 89 FLOGIs or F Disc per port for the 9100 series switches.

<sup>1</sup> For information about the 48-port module, see [“C-series Fibre Channel switches”](#) (page 115).

<sup>2</sup> MDS 9124, MDS 9124e, and MDS 9134 support a maximum of 16 VSANs.

<sup>3</sup> MDS NX-OS release 5.0(x) configuration limits

**NOTE:** Not all topologies can support the maximum port or switch count.

## ISL maximums

You can use all full-rate ports on all C-series Fibre Channel switches for ISLs, with a maximum of one half of the total ISL port count configured to the same destination switch. [Table 50](#) (page 128) lists the ISL maximums for switches with higher port counts.

**Table 50 ISL maximums**

Switch	Total number of available user ports	Number of ports allowed as ISLs
SN6000C	48	48 at 8 Gb/s
MDS 9xxx 48-port 8 Gb/s Host-Optimized Fibre Channel Switching Module	48	48 at 1 Gb/s
MDS 95xx 48-port 8 Gb/s Performance Fibre Channel Switching Module	48	48 at 2 Gb/s
MDS 95xx 24-port 8 Gb/s Fibre Channel Switching Module	24	24 at 4 Gb/s
MDS 9xxx 48-port 4 Gb/s Fibre Channel Switching Module	48	48 at 1 Gb/s
MDS 9xxx 24-port 4 Gb/s Fibre Channel Switching Module	24	24 at 2 Gb/s
MDS 9xxx 12-port 4 Gb/s Fibre Channel Switching Module	12	12 at 4 Gb/s
MDS 9020	20	20 at 4 Gb/s
MDS 9xxx 4-port 10 Gb/s Fibre Channel Switching Module	4	4 at 10 Gb/s
MDS 9xxx 32-port 2 Gb/s Fibre Channel Switching Module	32	8 at 2 Gb/s <sup>1</sup>
MDS 9xxx 16-port 2 Gb/s Fibre Channel Switching Module	16	16 at 2 Gb/s
MDS 9xxx 14/2 Multiprotocol Services Module	14	14 at 2 Gb/s
MDS 9xxx 18/4 Multiservice Module	18	18 at 2 Gb/s



**Table 50 ISL maximums (continued)**

Switch	Total number of available user ports	Number of ports allowed as ISLs
MDS 9120	20	4 at 2 Gb/s <sup>1</sup>
MDS 9140	40	8 at 2 Gb/s <sup>1</sup>
MDS 9124 Fabric	24	24 at 4 Gb/s
MDS 9124e Fabric	8	8 at 4 Gb/s
MDS 9134 Fabric	34	32 at 4 Gb/s 2 at 10 Gb/s

<sup>1</sup> Three user ports must be disabled for each port configured as an ISL.

## Zoning limits and enforcement

Table 51 (page 129) lists the zoning limits for C-series Fibre Channel switches. See also “Zoning” (page 376).

**Table 51 Zoning limits for C-series Fibre Channel switches**

Rule number	Description <sup>1</sup>
1	Maximum number of zones for a fabric with VSANs is 8,000.
2	Maximum number of zone members for a fabric with VSANs is 20,000.

<sup>1</sup> VSANs and IVR are not supported on the MDS 9020. IVR is not supported on the MDS 9124, MDS 9124e, and MDS 9134.

Table 52 (page 129) describes zoning enforcement for C-series Fibre Channel switches.

**Table 52 Zoning enforcement for C-series Fibre Channel switches**

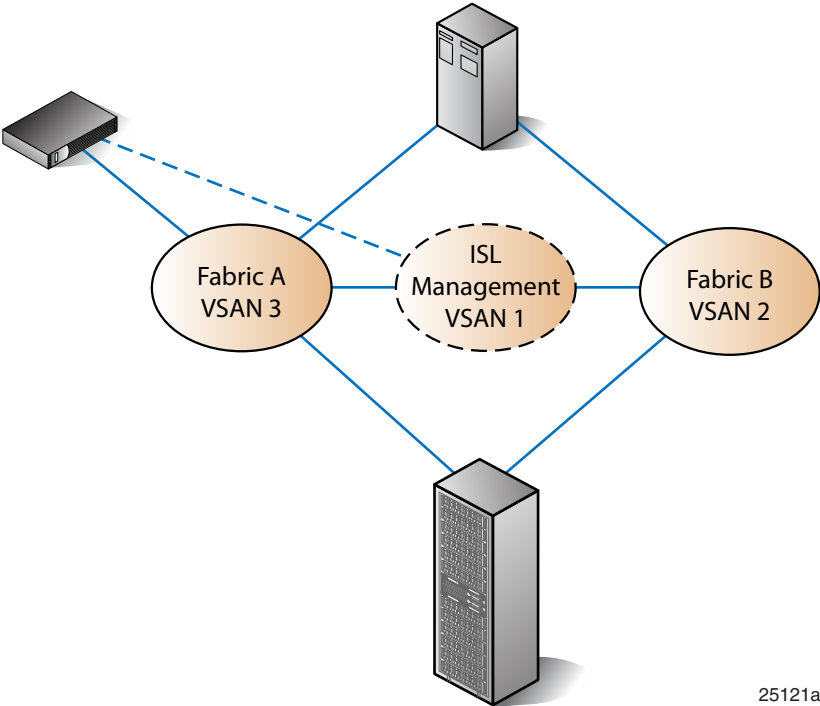
Switch	Configuration	Enforcement	Comments
MDS 8Gb fabric switch for HP BladeSystem c-Class SN6000C MDS 9506 MDS 9509 MDS 9513 MDS 9216A/9216i MDS 9222i MDS 9120 MDS 9140 MDS 9124 MDS 9124e MDS 9134	Defines zones using Domain ID and port number Defines zones using WWNs only Defines zones using Domain ID and port number with WWNs	Access authorization at frame level in hardware	Hard zoning
MDS 9020	pWWN FC alias	Name Server directory-based authentication	Soft zoning

## C-series VSAN high availability

Figure 53 (page 130) shows a typical high-availability configuration with server and storage connections to different fabrics. It provides two paths for data access between servers and storage. The addition of an ISL between Fabric A and Fabric B enables you to manage both fabrics through

a single management server. This configuration is classified as a level 3 high-availability configuration. For information about SAN data availability levels, see “Data availability” (page 38).

**Figure 53 C-series Fibre Channel switch high-availability VSAN management configuration**



25121a

---

## 7 H-series switches and fabric rules

This chapter describes the fabric rules for the HP H-series switches. It describes the following topics:

- “H-series switches” (page 131)
- “Fabric rules” (page 134)

### H-series switches

The SN6000 Fibre Channel Switches:

- Can be used as core or edge switches. When configured in a core-edge fabric topology, a core switch typically connects to other switches in the SAN; an edge switch typically connects to servers and storage.
- Support multi-switch networks with up to 30 switches and 7 hops.
- Have eight or twenty (all) base 8 Gb ports and four stacking (high speed ISL) 10 Gb ports. For the switch model with eight 8 Gb ports enabled, you can upgrade to twenty 8 Gb ports in 4-port increments. You can upgrade the 10 Gb ports to 20 Gb. Using the 10 Gb/20 Gb ports for ISLs reduces or eliminates the need to use 8 Gb device ports for ISLs, allowing more device connections to the switch.
- Can make use of the 10 Gb/20 Gb stacking ports to configure up to six switches into a stack (single manageable unit), resulting in up to 120 usable device ports in a stack.
- Support plug-and-play compatibility.
- Have two hot-swappable power supplies or one replaceable power supply depending on model.

The 8/20q Fibre Channel Switches:

- Can be used as core or edge switches. When configured in a core-edge fabric topology, a *core switch* typically connects to other switches in the SAN; an *edge switch* typically connects to servers and storage.
- Support multi-switch networks with up to 30 switches and 7 hops.
- Have 8 or 16 base ports, which you can upgrade to 20 ports in 4-port increments.
- Support plug-and-play compatibility.

The 8 Gb Simple SAN Connection Kit offers:

- Ease of use
- Cost efficiency

The kit includes:

- An 8/20q Fibre Channel Switch (8 base ports that can be upgraded to 12, 16, or 20)
  - (4) 81Q PCI-e FC HBAs
  - (10) 8-Gb SFP+ transceivers (SFP+ is the expanded standard for 8-Gb Fibre Channel support.)
  - (6) Fiber optic cables
- Scalability

You can expand the SAN in four port increments using license keys and additional HBAs and switches.

## Model numbering

The SN6000 Fibre Channel Switch uses the numbering scheme of SN6xxx, which indicates a mid-range size switch type.

The 8/20q Fibre Channel Switch uses the numbering scheme  $x/y$ :

- $x$ —The highest speed at which the switch ports can operate, measured in Gb/s
- $y$ —The total number of switch ports available

For example, the 8/20q Fibre Channel Switch is an 8 Gb/s switch with up to 20 ports.

## Model naming

The SN6000 Fibre Channel Switch is available as a standalone switch. The 8/20q Fibre Channel Switch is available as a standalone switch or as part of the 8 Gb Simple SAN Connection Kit. The 8 Gb Simple SAN Connection Kit includes the 8/20q Fibre Channel Switch with eight ports enabled, four HBAs, ten 8 Gb SFP+ transceivers, and six 5-meter fiber optic cables. Both switches and kit include comprehensive management software—the HP StorageWorks Simple SAN Connection Manager (SSCM).

## Switch models

HP supports all H-series switches in a fabric if:

- All H-series switches in a single-fabric or multi-fabric SAN use the appropriate firmware version (see [Table 53 \(page 132\)](#)). When updating switch firmware, you can use two successive switch firmware versions temporarily.
- The fabric rules are followed. See “Fabric rules” ([page 134](#)).

For the latest information on supported H-series switch and firmware versions, see one of the HP SAN Infrastructure websites. For the SN6000 Fibre Channel Switch, see <http://www.hp.com/go/SN6000>; for the 8/20q Fibre Channel Switch, see <http://www.hp.com/go/8Gb-SSC>.

**Table 53 H-series switches**

HP switch	Firmware version	Number of ports
HP StorageWorks SN6000 Fibre Channel Switch	8.0.4.04.00	8/12/16/20/24
HP StorageWorks 8/20q Fibre Channel Switch	8.0.4.04.00	8/12/16/20
	8.0.2.02.00 7.6.0.02.00	

## Features

Features of the SN6000 Fibre Channel Switch include:

- Four stacking (high speed ISL) 10 Gb/s ports, upgradable to 20 Gb/s. Using the stacking ports for ISLs enables configuring up to six switches into a stack (single manageable unit). For more information, see the *HP StorageWorks SN6000 Fibre Channel Switch Installation and Reference Guide*, available at the HP SAN Infrastructure website <http://www.hp.com/go/SN6000>.
- Autodiscovering, autonegotiating, and self-configuring ports (2/4/8 Gb/s and 10/20 Gb/s)
- Nonblocking, full-bandwidth architecture
- Adaptive trunking
- Advanced security features (CHAP, RADIUS authentication, SSL/SSH, and port binding)

- TR\_Port configuration, which enables connectivity of SN6000 Fibre Channel Switches to B-series or C-series fabrics
- High-availability features:
  - Two hot-swappable power supplies in dual power supply model
  - Hot-swappable 8 Gb SFP+ and 10 Gb XPAK optical transceivers
  - NDCLA
  - Nondisruptive port-license activation
  - Support in high-availability redundant configurations

Features of the 8/20q Fibre Channel Switch include:

- Autodiscovering, autonegotiating, and self-configuring ports (2/4/8 Gb/s)
- Nonblocking, full-bandwidth architecture
- Adaptive trunking
- Advanced security features (CHAP, RADIUS authentication, SSL/SSH, and port binding)
- TR\_Port configuration, which enables connectivity of 8/20q Fibre Channel Switches to B-series or C-series fabrics
- High-availability features:
  - Hot-swappable 8 Gb SFP+ optical transceivers
  - NDCLA
  - Nondisruptive port-license activation
  - Support in high-availability redundant configurations

Features of the 8 Gb Simple SAN Connection Kit include:

- Includes all components required to create a four-host SAN infrastructure with connection to a single- or dual-controller storage target—HBA, switch, SFP+ optical transceivers, and cables included in one SKU
- Can be installed by HP, a VAR, or the customer
- Provides end-to-end management, including switch, HBA, and storage provisioning

Features of the SSCM include:

- Comprehensive, easy-to-use management software
- Simple wizard-based installation
- One GUI for the H-series switches, HBAs, and MSA/EVA storage systems
- Configuration, zoning, and application-based MSA/EVA storage setup and provisioning
- Topology mapping
- HBA and switch device management
- Automated firmware and HBA software update notifications and downloads

Enterprise Fabric Management Suite (EFMS) is a separately licensed, workstation-based fabric management GUI. Features of EFMS include:

- All the capabilities of QuickTools
- Fabric tracker for monitoring firmware versions
- Port threshold alarm configuration
- Performance view for port performance monitoring

- Extended credits wizard to increase port distance capabilities (see [Table 78 \(page 150\)](#) and [“H-series switch settings” \(page 257\)](#))
- mPort technology for moveable port licenses
- Media diagnostics
- Fibre Channel tracing and connection verification

You can download EFMS with a 30-day trial license from the HP website [www.hp.com/go/EFMS](http://www.hp.com/go/EFMS). For information about configuring the switch using EFMS, see the *HP StorageWorks 8/20q and SN6000 Fibre Channel Switch Enterprise Fabric Management Suite User Guide*.

[Table 54 \(page 134\)](#) describes the high-availability features of the H-series switches.

**Table 54 H-series switches high-availability features**

Model	Redundant/ hot-swappable power	Redundant/ hot-swappable cooling	Nondisruptive code load/ activation	Nondisruptive port expansion/upgrade
HP StorageWorks SN6000 Stackable Dual Power Supply Fibre Channel Switch	Yes/Yes	Yes/Yes	Yes/Yes	Yes
HP StorageWorks SN6000 Stackable Single Power Supply Fibre Channel Switch	No/No	No/No	Yes/Yes	Yes
HP StorageWorks 8/20q Fibre Channel Switch	No/No	No/No	Yes/Yes	Yes

## Usage

The H-series switches can be used as core or edge switches in fabrics with up to 600 ports (including 8 Gb ISLs, but not including 10 Gb/20 Gb stacking port ISLs).

The TR feature allows you to connect H-series switches to B-series or C-series fabrics, enabling the sharing of server and storage resources between fabrics through the industry-standard NPIV protocol.

## Fabric rules

This section describes the SAN fabric rules that apply to all H-series switch SANs.

When using H-series switches in a fabric:

- Use the HP default switch settings for all configurations.
- All switch and fabric rules apply to fabrics implemented with the switch firmware versions listed in [Table 53 \(page 132\)](#).

## Servers, operating systems, and storage products

The fabric rules for H-series switches apply to SANs that include the servers, operating systems, and storage products listed in [Table 55 \(page 135\)](#).

**Table 55 H-series switches servers, operating systems, and storage system support**

Servers	Operating systems	Storage products
<ul style="list-style-type: none"> <li>HP ProLiant ML-Series and DL-Series servers</li> <li>HP Integrity servers</li> <li>HP BladeSystem c-Class ProLiant Server Blades</li> <li>HP BladeSystem c-Class Integrity Server Blades</li> </ul>	<ul style="list-style-type: none"> <li>Microsoft Windows</li> <li>VMware ESX</li> <li>Red Hat Enterprise Linux (RHEL)</li> <li>SUSE Linux Enterprise Server (SLES)</li> <li>HP-UX</li> <li>Citrix Xen Enterprise</li> <li>Integrated Citrix XenServer</li> </ul>	<ul style="list-style-type: none"> <li>P6300/P6500 EVA<sup>1</sup></li> <li>EVA8400/6400</li> <li>EVA4400</li> <li>EVA8100/8000/6100/6000/4100/4000</li> <li>MSA2000fc G2 (MSA2300fc)/MSA2000fc</li> <li>MSA1000/1500 (active/active only)</li> <li>P9500</li> <li>P2000 G3 FC MSA</li> <li>XP24000/20000, XP12000/10000</li> <li>3PAR F-Class, T-Class</li> <li>P6000 Continuous Access<sup>2</sup></li> <li>EBS products: See the EBS compatibility matrix at <a href="http://www.hp.com/go/ebs">http://www.hp.com/go/ebs</a>.</li> </ul>

<sup>1</sup> SSCM does not support P6300/P6500.

<sup>2</sup> P6000 Continuous Access is supported with EVA4400/6400/8400 and firmware XCS version 09534000 with the data replication protocol option set for "HP SCSI FC Compliant Data Replication Protocol," and with switch firmware 8.0.4.04.00.

For current storage system support and firmware, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

For more configuration information, see the *HP P6000 Continuous Access Implementation Guide*.

The operating systems listed in [Table 55 \(page 135\)](#) may not be supported with every storage system listed. For current operating system and storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

See "Heterogeneous server rules" (page 159) and the following chapters:

- "P2000 and MSA storage system rules" (page 203)
- "P6000/EVA storage system rules" (page 213)

To support SSCM, EFMS, and other HP software, the server running the SSCM management software (the management station) must meet the following minimum requirements:

- Microsoft Windows Server 2003 R2 x64/x86 with SP 2 and Microsoft hotfix QFE932755, available on the Microsoft website at <http://support.microsoft.com/kb/932755>
- Microsoft Windows Storage Server 2003 R2 x64/x86 with SP 2 and Microsoft hotfix QFE932755, available on the Microsoft website at <http://support.microsoft.com/kb/932755>
- Microsoft Windows Server 2008 x64/x86 with Service Pack 2
- Microsoft Windows Server 2008 R2 x64/x86

## Fabric rules for H-series switches

[Table 56 \(page 136\)](#) describes the rules for creating a SAN with H-series switches.

**Table 56 H-series switches fabric rules**

Rule number	Description
1	Supports up to 30 switches with up to 600 total 8 Gb/s ports, up to 120 10 Gb/20 Gb stacking ports, and up to 512 user ports in a fabric.
2	Supports a maximum of 7 switch hops (8 switches) between any 2 communicating devices.
3	8/20q Fibre Channel Switch supports from 8 to 20 8 Gb/s ports, in 4-port increments; SN6000 Fibre Channel Switch supports 8 to 20 8 Gb/s ports, in 4-port increments, plus 4 stacking ports.

## ISL maximums

When designing a fabric using 8-port, 12-port, 16-port, or 20-port switches, you can use up to six 8 Gb ports per switch as ISLs.

## Fabric rules for H-series switches with TR

This section describes the fabric rules for H-series switches with TR connecting to B-series or C-series fabrics.

The TR feature provides inter-fabric routing on a per-port basis, allowing controlled access between devices on an H-series switch (local) fabric and devices on a remote fabric consisting of B-series or C-series switches. The establishment of a routed connection using TR maintains a high level of isolation between fabrics. A transparent route between two devices consists of a connection from a TR\_Port on an H-series switch to a switch in the remote fabric, a mapping of the two devices to be routed together, and an IFZ for the routed devices in both fabrics.

Each fabric contains a matching IFZ and each IFZ contains three WWN members: the local device, the remote device, and the TR\_Port connected to the remote fabric. This inter-fabric connection uses the Fibre Channel industry-standard NPIV, making local and remote devices accessible to each other and maintaining local and remote fabrics as separate fabrics.

You can connect multiple H-series switches to one or more remote fabrics using multiple TR\_Ports. Local and remote devices are identified by their respective worldwide port names.

You can configure transparent routing using QuickTools, EFMS, or the CLI. HP recommends that you use QuickTools or EFMS because they validate your entries, manage the zone mapping for the local fabric, and create a list of zoning commands you can run in a script on a B-series or C-series SAN switch. For more information and important configuration details, see the following documents:

- *HP StorageWorks SN6000 Fibre Channel Switch QuickTools Switch Management User Guide*
- *HP StorageWorks SN6000 Fibre Channel Switch Command Line Interface Guide*
- *HP StorageWorks 8/20q Fibre Channel Switch QuickTools Switch Management User Guide*
- *HP StorageWorks 8/20q Fibre Channel Switch Command Line Interface Guide*
- *HP StorageWorks 8/20q and SN6000 Fibre Channel Switch Enterprise Fabric Management Suite User Guide*



Table 57 (page 137) describes the supported remote fabric switches you can use to connect to an H-series switch TR\_Port.

**Table 57 Supported switches in an H-series switch with TR remote fabric**

HP switch name	Firmware version
Brocade 8Gb SAN Switch for HP BladeSystem c-Class HP StorageWorks 8/8 Base, SAN Switch HP StorageWorks 8/24 SAN Switch HP StorageWorks 8/40 SAN Switch HP StorageWorks 8/80 SAN Switch HP StorageWorks EVA4400 Embedded Switch Module, 8 Gb Brocade Brocade 4Gb SAN Switch for HP c-Class BladeSystem HP StorageWorks 4/8 SAN Switch HP StorageWorks 4/16 SAN Switch HP StorageWorks SAN Switch 4/32 HP StorageWorks SAN Switch 4/32B HP StorageWorks SAN Switch 4/64 HP StorageWorks 4/256 SAN Director	6.2.1b (or later)
Cisco MDS 9513 Multilayer Director Switch Cisco MDS 9509 Multilayer Director Switch Cisco MDS 9506 Multilayer Director Switch Cisco MDS 9222i Fabric Switch Cisco MDS 9134 Fabric Switch Cisco MDS 9124 Fabric Switch Cisco MDS 9124e Fabric Switch for c-Class BladeSystem	5.0(1a) (or later)
Cisco MDS 9513 Multilayer Director Switch Cisco MDS 9509 Multilayer Director Switch Cisco MDS 9506 Multilayer Director Switch Cisco MDS 9222i Fabric Switch Cisco MDS 9134 Fabric Switch Cisco MDS 9124 Fabric Switch Cisco MDS 9124e Fabric Switch for c-Class BladeSystem	4.2(1a) (or later)
Cisco MDS 9509 Multilayer Director Switch Cisco MDS 9506 Multilayer Director Switch Cisco MDS 9216A Multilayer Fabric Switch	3.3(4) (or later)

Table 58 (page 137) describes the rules for creating a SAN with an H-series switch with TR and remote B-series or C-series fabrics.

**Table 58 H-series switch TR and remote B-series or C-series fabric rules**

Rule number	Description
1	The local fabric can consist of one or more H-series switches connected by ISLs.
2	Any H-series switch 8 Gb port can be configured as a TR_Port.
3	The switch port in the remote fabric to which the TR_Port is connected must be enabled for NPIV support.
4	A TR_Port can support a maximum of 32 local device to remote device mappings.
5	A local device can be mapped to one or more devices on only one remote fabric.

**Table 58 H-series switch TR and remote B-series or C-series fabric rules** *(continued)*

Rule number	Description
6	Local devices on an H-series switch can be mapped to different remote fabrics.
7	For mappings between an H-series switch and a remote fabric, each local device or remote device can be mapped over one TR_Port. Additional mappings to either device must use that TR_Port.
8	Local devices connected to different local switches can be mapped to the same remote device over one TR_Port on each local switch.
9	A local device cannot be mapped over an E_Port to a local switch, then over a TR_Port to the remote device. The local switch to which the local device is connected must connect directly to the remote fabric over a TR_Port.
10	In order for local devices to discover remote devices, the corresponding IFZs must be activated on both the local and remote fabrics. To remove a mapping, you must remove the local IFZ and the corresponding remote IFZ.
11	When a local device is mapped over a TR_Port to a remote device, the local device and its TR_Port appear as an NPIV-connected device in the remote fabric. It is possible, though not recommended, to map the local device over a second TR_Port to a local device in a second local fabric. In this case, if you merge the two local fabrics, the transparent route becomes inactive for the devices that now have a path over an ISL, and an alarm is generated.

- ⓘ **IMPORTANT:** SSCM version 3.00 (or later) is required for the HP SN6000 Fibre Channel Switch. SSCM version 2.10 and later can manage the 8/20q Fibre Channel Switch, and version 3.00 (or later) can manage the SN6000 Fibre Channel Switch, with active transparent router ports (TR\_Ports). However, SSCM cannot manage or discover remote switches or devices in the remote fabric. The Physical Connection map displays the remote fabric as a disabled switch, indicating that SSCM cannot manage the switch. SSCM version 2.00 and earlier does not support the management of fabrics that include switches with active TR\_Ports, and may disrupt communication between the 8/20q Fibre Channel Switch or SN6000 Fibre Channel Switch and the remote fabric.

## Zoning limits and enforcement

Table 59 (page 138) lists zoning limits for H-series switches. See also “Zoning” (page 376).

**Table 59 H-series switch zoning limits**

Rule number	Description
1	Maximum number of zone sets is 256.
2	Maximum number of zones is 2,000.
3	Maximum number of zones in a zone set is 2,000.
4	Maximum number of zone members in a zone is 2,000.
5	Maximum number of zone members per fabric is 10,000.

Table 60 (page 139) describes zoning enforcement for H-series switches.

**Table 60 Zoning enforcement for H-series switches**

Switch	Configuration	Enforcement	Comments
HP StorageWorks SN6000 Fibre Channel Switch HP StorageWorks 8/20q Fibre Channel Switch	Define zones using domain number and port number	Access authorization at frame level in hardware	Hard zoning <sup>1</sup>
	Define zones using WWNs only	Access authorization at frame level in hardware	Hard zoning <sup>1</sup>
	Define zones using combination of domain/port numbers and WWNs	Name Server directory-based authentication	Soft zoning

<sup>1</sup> Zoning is hardware enforced only on a switch port or device that is a member of no more than 8 zones whose combined membership does not exceed 64. Zoning is hardware enforced on a switch port if the number of logged-in devices plus the number of zone members with access to that port is 64 or less. If these requirements are not met, the port implements soft zoning.

## 8 M-series switches and fabric rules

This chapter describes the M-series switches and the fabric rules for building M-series fabrics. It describes the following topics:

- “M-series switches” (page 140)
- “Fibre Channel fabric rules” (page 142)

**NOTE:** For information about using switches from the different series in the same SAN or the same fabric, see “SAN fabric connectivity and switch interoperability rules” (page 146).

### M-series switches

The following M-series products have reached End Of Support Life (ESOL) and are no longer supported:

- HP StorageWorks Director 2/140
- HP StorageWorks Director 2/64
- HP StorageWorks Edge Switch 2/32
- HP StorageWorks Edge Switch 2/24
- HP StorageWorks Edge Switch 2/12
- StorageWorks HA-Fabric Manager (HAFM)
- HA Fabric Manager Appliance
- HA Fabric Manager Appliance FRU
- HP HA Fabric Manager SW UG kit

### Model naming

The M-series embedded switch for HP p-Class BladeSystem servers is called the McDATA 4 Gb SAN Switch.

### Switch models

Table 61 (page 140) describes the M-series blade system switch. Table 62 (page 141) describes the McDATA switches supported by HP for use in connectivity to HP Storage.

HP supports all M-series switches in a fabric if you:

- Use the firmware versions listed in Table 61 (page 140) and Table 62 (page 141). All like switch models must use the same version when configured in a single-fabric or multi-fabric SAN. The firmware for the McDATA 4 Gb SAN Switch for HP p-Class BladeSystem must be compatible with the firmware used on HP StorageWorks Edge Switches and Directors. When updating switch firmware, you can use two successive switch firmware versions temporarily in a single-fabric or multi-fabric SAN.
- Follow the fabric rules. See “Fibre Channel fabric rules” (page 142).

For the latest information on supported M-series switches and firmware versions, see the HP SAN Infrastructure website: <http://h18006.www1.hp.com/storage/saninfrastructure.html>.

**Table 61 M-series switches**

HP switch	Firmware version	Number of ports
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	6.4.0.13.00	8 internal
	6.4.0.11.00	2 external

**Table 62 McDATA/Brocade switches supported with M-series switches**

McDATA/Brocade switch	Firmware version	Number of ports
Sphereon 4700/M4700 Fabric Switch	09.09.09	8 to 32
Sphereon 4400/M4400 Fabric Switch	09.09.08 09.09.06 09.09.05 09.07.02 <sup>1</sup>	4 to 16
Intrepid 10000/Mi 10K Director	09.09.08 09.09.07 09.09.06 09.09.05 09.07.02 <sup>1</sup>	32 to 256

<sup>1</sup> Firmware 09.07.02 is only supported for fabrics which include the McDATA 4 Gb SAN Switch for HP p-Class BladeSystem

## Features

Features of the Director switch includes:

- Nonblocking port density to minimize floor-space usage
- Low power consumption and heat generation
- Redundant, hot-swappable switching and processor logic cards
- Nondisruptive failover of redundant components
- Automatic health checks of redundant FRUs

Table 63 (page 142) compares the high-availability features of M-series switches.

**Table 63 M-series high-availability feature comparison**

Model	Size (1U = 1.75 inch)	Redundant control processor/switching	Redundant active components
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	N/A	No	No

## Usage

Table 64 (page 142) describes the use of M-series switches as director switches.

**Table 64 Using M-series switches as director switches**

Model	1–500 total ports	501–1000 total ports	1001–1632 total ports
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	Good	Not recommended	

Table 65 (page 142) describes the use of M-series switches as edge switches.

**Table 65 Using M-series switches as edge switches**

Model	1–500 total ports	501–1000 total ports	1000–1632 total ports
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	Good	Not recommended	

## Fibre Channel fabric rules

The following SAN fabric rules apply to all M-series SANs. They also apply, in general, to HP P9000 (XP) Continuous Access and HP P6000 Continuous Access configurations. However, additional rules apply to HP Continuous Access implementations. For more information, see the *HP P6000 Continuous Access Implementation Guide*.

When using M-series switches in a fabric:

- Use the HP default settings for all current HP switch models.
- All switch and fabric rules apply to fabrics implemented with the switch firmware versions listed in Table 61 (page 140) and Table 62 (page 141).

## Operating systems and storage products

The fabric rules for M-series switches apply to SANs that include the operating systems and storage products listed in Table 66 (page 143).

**Table 66 M-series Fibre Channel switch operating system and storage system support**

Operating systems	Storage products
<ul style="list-style-type: none"> <li>• HP-UX</li> <li>• OpenVMS</li> <li>• Tru64 UNIX</li> <li>• IBM AIX</li> <li>• Linux</li> <li>• Microsoft Windows</li> <li>• Sun Solaris</li> <li>• VMware ESX</li> </ul>	<ul style="list-style-type: none"> <li>• EVA8400/6400</li> <li>• EVA4400</li> <li>• EVA8100/8000/6100/6000/4100/4000</li> <li>• EVA50000/3000</li> <li>• MSA2000fc G2 (MSA2300fc)/MSA2000fc</li> <li>• MSA1500/1000</li> <li>• P9500<sup>1</sup></li> <li>• P2000 G3 FC</li> <li>• XP24000/20000</li> <li>• XP12000/10000</li> <li>• XP1024/128</li> <li>• XP512/48</li> <li>• HP P9000 (XP) Continuous Access</li> <li>• HP P6000 Continuous Access See “HP P6000 Continuous Access SAN integration” (page 220)</li> </ul>

<sup>1</sup> Intrepid 10000/Mi10K Director support only

For current operating system and storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

See “Heterogeneous server rules” (page 159) and the following storage system rules chapters:

- “P2000 and MSA storage system rules” (page 203)
- “P6000/EVA storage system rules” (page 213)
- “P9000/XP and VA storage system rules” (page 227)

## Fabric rules for M-series switches

Table 67 (page 143) describes the rules for creating a SAN with M-series switches.

**Table 67 M-series fabric rules**

Rule number	Description
1	A maximum of 24 switches, 1,632 total ports, and 1,500 user ports are supported in a single fabric. See Table 69 (page 144). 1 Gb/s, 2 Gb/s, and 4 Gb/s switches are supported in the fabric if you do not exceed the following limit: <ul style="list-style-type: none"> <li>• For a fabric that contains the McDATA 4 Gb SAN Switch for HP p-Class BladeSystem, a maximum of 24 switches and 600 user ports (firmware 6.4.x.xx.xx)</li> </ul>
2	The supported limit of interconnected switches in a single fabric is 31, provided that any edge switch in the fabric is deployed at the edge of the fabric.
3	A maximum of three switch hops (four switches) is supported between any two devices.
4	A ring fabric topology is supported with a maximum of seven M-series switches.
5	Assign a unique domain number (domain ID) and a unique WWN to each switch. The switch configuration parameters must be the same for each switch. Do not configure any switches with a domain ID of 8, which is reserved for HP-UX.
6	HP requires that all switches in a single-fabric SAN or multi-fabric SAN use the same firmware version for the same switch models. You can use two successive switch firmware versions temporarily in one fabric or multiple fabrics when updating switch firmware.

---

**NOTE:** Not all topologies can support the maximum port or switch count.

---

## ISL maximums

Table 68 (page 144) lists the maximums for switches with higher port counts.

**Table 68 ISL maximums**

Switch	Total number of available user ports	Number of E_Ports allowed
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	10	2

## Zoning limits and enforcement

Table 69 (page 144) lists zoning limits for M-series and McDATA switches. See also “Zoning” (page 376).

**Table 69 M-series and McDATA zoning limits**

Zoning parameter	Firmware 09.xx.xx	Firmware 6.4.x.xx.xx
	Maximum value with M-series switches	Maximum value with McDATA 4 Gb SAN Switch for HP p-Class BladeSystem
Number of zone members in a zone	4,096	4,096
Number of zones in a zone set	2,048	2,047
Number of unique zone members in a zone set	4,096	4,096
Total number of zone members in a zone set (where a zone member can be in multiple zones)	4,096	4,096
Characters per zoning name	64	64
Number of end ports	2,048	600
Number of devices supported (including loop devices)	2,048	600



Table 70 (page 145) describes zoning enforcement for M-series and McDATA switches.

**Table 70 Zoning enforcement for M-series and McDATA switches**

Switches	Configuration	Enforcement	Comments
<b>McDATA</b> Sphereon 4700 Sphereon 4400 <b>Brocade</b> M4700 Fabric Switch M4400 Fabric Switch	Define zones using domain number and port number (HP M-series Homogeneous Fabric mode or McDATA Fabric mode only) Define zones using WWNs only Define zones using combination of domain/port numbers and WWNs (HP M-series Homogeneous Fabric mode or McDATA Fabric mode only)	Discovery authentication Name Server directory-based authentication Access authorization at frame level in hardware	Soft zoning or hard zoning (5.01.00-24 and later)
McDATA 4 Gb SAN Switch for HP p-Class BladeSystem	Define zones using domain number, port number (must be in McDATA Fabric mode if connected to an M-series fabric) Define zones using WWNs only Define zones using combination of domain/port numbers and WWNs (must be in McDATA Fabric mode if connected to an M-series fabric)	Discovery authentication Name Server directory-based authentication	Soft zoning

---

## 9 SAN fabric connectivity and switch interoperability rules

This chapter describes SAN fabric connectivity and interoperability rules. It describes the following topics:

- “SAN fabric connectivity rules” (page 146)
- “SAN fabric switch interoperability rules” (page 154)
- “Third-party switch support” (page 156)
- “SAN performance considerations” (page 156)

### SAN fabric connectivity rules

This section describes SAN fabric connectivity port interfaces, cables, and rules.

#### Switch port interfaces

The switch port interfaces are as follows:

- **E\_Port**—Provides switch-to-switch connectivity for ISLs
- **EX\_Port**—Connects a Fibre Channel router to an edge fabric
- **F\_Port**—Provides fabric-attached device connectivity for initiators (HBAs) and targets (storage ports)
- **FL\_Port**—Provides fabric-aware public loop connectivity with 24-bit Fibre Channel addressing capability
- **FCAL\_Port**—Provides private loop connectivity for 8-bit Fibre Channel addressable devices; (requires B-series QuickLoop feature)

QuickLoop enables private FC-AL initiators and targets to communicate through the switch. A target that is not configured with QuickLoop cannot communicate with a QuickLoop initiator. QuickLoop is supported only for specific B-series switches and legacy storage systems. For more information, contact an HP storage representative.

- **VE\_Port**—Virtual E\_Port used with a GbE port for FCIP tunneling
- **VEX\_Port**—Virtual EX\_Port used for Fibre Channel routing over a VE\_Port

Certain HP B-series (NPIV), C-series (NPV), and all H-series (NPIV) switches use an N\_Port ID Virtualization interface. NPIV and NPV are industry-standard Fibre Channel protocols that support the assignment of multiple Fibre Channel N\_Port addresses on the same physical link. See “Blade servers” (page 161).

#### Device port interfaces

The server HBA or storage system controller device node port interfaces are as follows:

- **N\_Port**—Connects a Fibre Channel device to a switch
- **NL\_Port**—Connects a Fibre Channel device to a Fibre Channel arbitrated loop or switch

#### Fiber optic cables

All 8 Gb/s, 4 Gb/s, and 2 Gb/s components use industry-standard LC connectors for fiber optic cable connections; all 1 Gb/s components use industry-standard SC connectors. Cables and adapters are available with SC connectors on one end and LC connectors on the other end.

Some fiber optic cable installations require SC or LC duplex couplers to couple the cable connector ends (for example, if you use wall jacks or connect to existing installed cables). HP supports the

use of duplex couplers, provided that you do not exceed the overall cable loss budget for that cable segment.

HP recommends the use of HP PremierFlex OM4 and OM3+ fiber optic cables for 50 micron cable installations. HP PremierFlex OM3+ fiber optic cables provide higher bend performance and improved signal-transmission integrity, providing significant improvements in signal quality over industry-standard OM3 fiber optic cable technologies.

For more information on PremierFlex, see:

- PremierFlex video: [http://h30428.www3.hp.com/?fr\\_story=12f70d1056643c69d0754fe491bd218c3859a2e0&rf=bn](http://h30428.www3.hp.com/?fr_story=12f70d1056643c69d0754fe491bd218c3859a2e0&rf=bn)
- PremierFlex data sheet: <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-6479ENUC.pdf>
- Cables portfolio data sheet: <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-4964ENUC.pdf>

**NOTE:** OM3+ is an HP designation indicating that the product exceeds all OM3 standards. OM3+ fiber optic cables specify a modal bandwidth of 3,000 MHz-km at 850 nm.

HP also supports industry standard fiber optic cable types OM4, OM3, OM2, and OM1. See “Fiber optic cable loss budgets” (page 148).

For cable part numbers, see the Fibre Channel switch QuickSpecs at <http://www.hp.com/go/QuickSpecs>.

Table 71 (page 147) describes the rules for fiber optic cable connections.

**Table 71 Rules for fiber optic cable connections**

Rule number	Description
1	Minimum bend radius—HP PremierFlex 50 micron OM4 and OM3+ fiber optic cable, 7 mm. Industry standard OM3, OM2, OM1, 25 mm for 50, 62.5, and 9 micron fiber optic cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber connections. The 62.5 micron fiber optic cable is acceptable for existing installations. See Table 75 (page 149) for supported maximum cable distances.
2	HP does not support a mixture of 50, 62.5, and 9 micron cable in the same cable segment. <b>Note:</b> A cable segment is defined as one or more cables connected serially between two transceivers.
3	HP does not support a mixture of different OM fiber cable types in the same cable segment.
4	For all cable segments, you must use the same transceiver type at both ends of the cable (for example, short-wave transceiver to short-wave transceiver, long-wave transceiver to long-wave transceiver). HP does not support mixing transceiver types in the same cable segment.
5	For fabric attachment, the minimum cable segment length between Fibre Channel components (transmitter and receiver) is: <ul style="list-style-type: none"> <li>• 0.5 m for 50 and 62.5 micron cable</li> <li>• 2.0 m for 9 micron cable</li> </ul> The minimum length does not apply to patch cords through a passive patch panel; it applies only to the total distance between the transmitter and receiver of each device connected through the patch panel.
6	For fiber optic cable lengths greater than 50 m, contact a third-party vendor. 50 micron cable must be duplex, tight-buffered multi-mode 50/125 $\mu$ m (Belcore GR-409 compliant). The connectors must be SC or LC duplex low metal (Belcore and IEC compliant). 9 micron cable must be duplex, tight-buffered, single-mode 9/125 $\mu$ m (Belcore GR-409 compliant). The connectors must be SC or LC duplex low metal (NTT-SC Belcore 326, IEC-874-19 SC compliant).

## Fiber optic cable loss budgets

Cable loss budgets are determined by the Fibre Channel Physical Interface Specification. See the standards at <http://www.incits.org/>. The maximum supported distances are based on modal bandwidth and type of fiber optic cable used. The different modal bandwidth and cable types that HP supports for Fibre Channel are:

- 62.5 micron fiber optic cable, modal bandwidth of 200 MHz-km at 850 nm (type OM1)
- 50 micron fiber optic cable, modal bandwidth of 500 MHz-km at 850 nm (type OM2)
- 50 micron fiber optic cable, modal bandwidth 2000 MHz-km at 850 nm (type OM3)
- 50 micron fiber optic cable, modal bandwidth 3000 MHz-km at 850 nm (type HP PremierFlex OM3+)

See “Fiber optic cables” (page 146) for more information on HP PremierFlex.

- 50 micron fiber optic cable, modal bandwidth 4700 MHz-km at 850 nm (type OM4, type HP PremierFlex OM4)

Table 73 (page 149) through Table 77 (page 150) list the maximum loss budgets for different interconnect speeds at specific distances. Table 78 (page 150) through Table 82 (page 153) and Table 85 (page 154) and Table 86 (page 154) list the maximum supported Fibre Channel distances based on switch-to-switch (ISL) or device-to-switch connectivity.

**NOTE:** The following tables do not specify media losses due to variances between different fiber optic cable manufacturers. In all cases, the maximum loss budget is the total channel insertion loss, which includes media losses based on the indicated fiber optic cable bandwidth. Table 83 (page 153) and Table 84 (page 153) list the maximum supported ATM and FCIP delays.

HP supports the use of optical fiber patch panels. The total channel insertion loss between the transmitter and receiver for the cable segment routed through the patch panel must not exceed the maximum listed for the connector and cable type.

**NOTE:** Channel insertion loss is the combined passive loss from connectors, splices, and media between the transmitter and receiver.

A mated connector pair is defined as a device or switch transceiver-to-cable connection, or a cable-to-cable connection when using a passive coupler in a patch panel. A typical (nonpatch panel) installation has two mated pairs per cable segment, one for each end of the cable. Additional mated connector pairs are allowed between two transceivers, provided the total loss of all connectors and cables does not exceed the total channel insertion loss.

Table 73 (page 149) lists the 10 Gb/s fiber optic cable loss budgets when using OM1, OM2, OM3, or OM3+ multi-mode fiber optic cable, and single-mode fiber optic cable.

Table 72 (page 148) lists the 16 Gb/s fiber optic cable loss budgets when using OM2, OM3, OM3+, or OM4 multi-mode fiber optic cable, and single-mode fiber optic cable.

### Table 72 16 Gb/s Fibre Channel fiber optic cable loss budgets

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
50/125 micron (OM2 500 MHz-km at 850 nm)	35 m	1.68 dB	0.75 dB
50/125 micron (OM3 2000 MHz-km at 850 nm, OM3+ 3000 MHz-km at 850 nm)	100 m	1.86 dB	
50/125 micron (OM4 4700 MHz-km at 850 nm)	125 m	1.95 dB	

**Table 73 10 Gb/s Fibre Channel fiber optic cable loss budgets**

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
62.5/125 micron (OM1 200 MHz-km at 850 nm)	33 m	1.6 dB	0.75 dB
50/125 micron (OM2 500 MHz-km at 850 nm)	82 m	1.8 dB	
50/125 micron (OM3 2000 MHz-km at 850 nm, OM3+ 3000 MHz-km at 850 nm)	300 m	2.6 dB	
9/125 micron (singlemode)	10 km	6.0 dB	

Table 74 (page 149) lists the 8 Gb/s fiber optic cable loss budgets when using OM1, OM2, OM3, OM3+, or OM4 multi-mode fiber optic cable, and single-mode fiber optic cable.

**Table 74 8 Gb/s Fibre Channel fiber optic cable loss budgets**

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
62.5/125 micron (OM1 200 MHz-km at 850 nm)	21 m	1.58 dB	0.75 dB
50/125 micron (OM2 500 MHz-km at 850 nm)	50 m	1.68 dB	
50/125 micron (OM3 2000 MHz-km at 850 nm, OM3+ 3000 MHz-km at 850 nm)	150 m	2.04 dB	
50/125 micron (OM4 4700 MHz-km at 850 nm)	190 m	2.19 dB	
9/125 micron (single-mode)	10 km	6.4 dB <sup>1</sup>	

<sup>1</sup> This number assumes use of low-loss fiber optic cables for 8 Gb/s speeds.

Table 75 (page 149) lists the 4 Gb/s fiber optic cable loss budgets when using OM1, OM2, OM3, OM3+, or OM4 multi-mode fiber optic cable, and single-mode fiber optic cable.

**Table 75 4 Gb/s Fibre Channel fiber optic cable loss budgets (nominal bandwidth)**

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
62.5/125 micron (OM1 200 MHz-km at 850 nm)	70 m	1.78 dB	0.75 dB
50/125 micron (OM2 500 MHz-km at 850 nm)	150 m	2.06 dB	
50/125 micron (OM3 2000 MHz-km at 850 nm, OM3+ 3000 MHz-km at 850 nm)	380 m	2.48 dB	
50/125 micron (OM4 4700 MHz-km at 850 nm)	400 m	2.95 dB	
9/125 micron (single-mode)	10 km	7.80 dB	
	30 km	19 dB	

Table 76 (page 150) lists the 2 Gb/s fiber optic cable loss budgets when using OM1, OM2, or OM3 multi-mode fiber optic cable, and single-mode fiber optic cable.

**Table 76 2 Gb/s Fibre Channel fiber optic cable loss budgets**

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
62.5/125 micron (OM1 200 MHz-km at 850 nm)	150 m	2.1 dB	0.75 dB
50/125 micron (OM2 500 MHz-km at 850 nm)	300 m	2.62 dB	
50/125 micron (OM3 2000 MHz-km at 850 nm)	500 m	3.31 dB	
9/125 micron (single-mode)	10 km	7.8 dB	
	35 km	21.5 dB	

Table 77 (page 150) lists the 1 Gb/s fiber optic cable loss budgets when using nominal bandwidth OM1 or OM2 multi-mode fiber optic cable, and single-mode fiber optic cable.

**Table 77 1 Gb/s Fibre Channel fiber optic cable loss budgets (nominal bandwidth)**

Cable	Maximum distance per cable segment	Total channel insertion loss	Maximum loss per mated connector pair
62.5/125 micron (OM1 200 MHz-km at 850 nm)	200 m	3.0 dB	0.75 dB
50/125 micron (OM2 500 MHz-km at 850 nm)	500 m	3.85 dB	
9/125 micron (single-mode)	10 km	7.8 dB	
	35 km	21.5 dB	
	100 km	19 dB	

## Storage product interface, switches, and transport distance rules

Table 78 (page 150) through Table 82 (page 153) and Table 85 (page 154) and Table 86 (page 154) describe the maximum distances supported for each cable segment type (switch-to-switch or device-to-switch) for each interface and transport type. Unless otherwise specified, the distances specified apply to both switch-to-switch (ISL) connectivity and device-to-switch connectivity.

Table 78 (page 150) describes the distance rules for 16 Gb/s or 8 Gb/s Fibre Channel connections when using 16 Gb/s or 8 Gb/s Fibre Channel switch models.

**Table 78 Fibre Channel distance rules for 16 Gb/s and 8 Gb/s switch models (B-series, C-series, and H-series)**

Interface/transport	Supported distances			Supported storage products
50 micron multi-mode fiber optic cable and short-wave SFP+ transceivers	<b>OM2 fiber</b>	<b>OM3, OM3+ fiber</b>	<b>OM4 fiber</b>	Heterogeneous SAN servers, Fibre Channel switches, and storage systems HP P6000 Continuous Access (B-series, C-series, and H-series only)
	35 m at 16 Gb/s 50 m at 8 Gb/s 150 m at 4 Gb/s 300 m at 2 Gb/s	100 m at 16 Gb/s 150 m at 8 Gb/s 380 m at 4 Gb/s 500 m at 2 Gb/s	125 m at 16 Gb/s 190 m at 8 Gb/s 400 m at 4 G/s	
62.5 micron <sup>1</sup> multi-mode fiber optic	<b>OM1 fiber</b>			HP P9000 (XP) Continuous Access
	21 m ISL at 8 Gb/s			

**Table 78 Fibre Channel distance rules for 16 Gb/s and 8 Gb/s switch models (B-series, C-series, and H-series)** (continued)

Interface/transport	Supported distances	Supported storage products
cable and short-wave SFP+ transceivers		(B-series and C-series only)
9 micron single-mode fiber optic cable and long-wave SFPs	10 km ISL at 8 Gb/s (H-series with additional buffer credits allocated) <sup>2</sup> 10 km ISL at 8 Gb/s (B-series and C-series only) 25 km ISL at 8 Gb/s (B-series only) 3.3 km ISL at 8 Gb/s (H-series only, base switch) 6.6 km ISL at 4 Gb/s (H-series only, base switch) 10 km ISL at 2 Gb/s (H-series only, base switch)	
Fibre Channel using WDM <sup>3</sup>	40 km ISL at 8 Gb/s (B-series only) 100 km ISL at 4 Gb/s 250 km at 2 Gb/s	

<sup>1</sup> Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

<sup>2</sup> You can use EFMS to allocate more buffer credits to ports of an H-series switch to achieve increased distance up to the limit of the SFP capability.

<sup>3</sup> WDM distance is the maximum distance for the WDM link.

Table 79 (page 151) describes the distance rules for 10 Gb/s Fibre Channel ISL connections when using 8 Gb/s Fibre Channel switch models with 10 Gb/s ISL ports.

**Table 79 Distance rules for 8 Gb switch models with 10 Gb/s Fibre Channel ISL ports (B-series, C-series, and H-series switches)**

Interface/transport	Supported distances	
CX4 Copper	15 m ISL at 10 Gb/s (C-series only) 20 m ISL at 20 Gb/s (H-series only)	
50 micron multi-mode fiber optic cable and short-wave SFP+ transceivers	<b>OM2 fiber</b> 82 m ISL at 10 Gb/s	<b>OM3, OM3+ fiber</b> 300 m ISL at 10 Gb/s
62.5 micron <sup>1</sup> multi-mode fiber optic cable and short-wave SFP+ transceivers	33 m ISL at 10 Gb/s (H-series only)	
9 micron single-mode fiber optic cable and long-wave SFPs	10 km ISL at 10 Gb/s (B-series and C-series only) 40 km ISL at 10 Gb/s (C-series only)	

<sup>1</sup> Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

Table 80 (page 151) describes the distance rules for 4 Gb/s Fibre Channel connections when using 4 Gb/s Fibre Channel switch models.

**Table 80 Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches)**

Interface/transport	Supported distances			Supported storage products
50 micron multi-mode fiber optic cable and short-wave SFPs	<b>OM2 fiber</b> 150 m at 4 Gb/s 300 m at 2 Gb/s	<b>OM3, OM3+ fiber</b> 380 m at 4 Gb/s 500 m at 2 Gb/s	<b>OM4 fiber</b> 400 m at 4 G/s	Heterogeneous SAN servers, Fibre Channel switches, and storage systems

**Table 80 Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches) (continued)**

Interface/transport	Supported distances		Supported storage products
	500 m at 1 Gb/s		
62.5 micron <sup>1</sup> multi-mode fiber optic cable and short-wave SFPs	<b>OM1 fiber</b> 70 m ISL at 4 Gb/s		HP P6000 Continuous Access and HP P9000 (XP) Continuous Access  EBS
9 micron single-mode fiber optic cable and long-wave SFPs	4 km ISL at 4 Gb/s (C-series only) 10 km ISL at 4 Gb/s		
9 micron single-mode fiber optic cable and extended-reach SFPs	30 km ISL at 4 Gb/s (B-series 8 Gb/s switches only) 35 km ISL at 2 Gb/s with 2 Gb/s SFP (B-series only)		
Fibre Channel using WDM <sup>2</sup>	100 km at 4 Gb/s 250 km at 2 Gb/s 500 km at 1 Gb/s		

<sup>1</sup> Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

<sup>2</sup> WDM distance is the maximum distance for the WDM link. For B-series switches, these distances are supported with firmware 5.x (or later).

Table 81 (page 152) describes the distance rules for 2 Gb/s Fibre Channel connections when using 2 Gb/s Fibre Channel switch models.

**Table 81 Fibre Channel distance rules for 2 Gb/s switch models (B-series, C-series, and M-series switches)**

Interface/transport	Supported distances		Supported storage products
50 micron multi-mode fiber optic cable and short-wave SFPs	<b>OM2 fiber</b> 300 m at 2 Gb/s 500 m at 1 Gb/s	<b>OM3, OM3+ fiber</b> 500 m at 2 Gb/s	Heterogeneous SAN servers, Fibre Channel switches, and storage systems  HP P6000 Continuous Access and HP P9000 (XP) Continuous Access  EBS
62.5 micron <sup>1</sup> multi-mode fiber optic cable and short-wave SFPs	<b>OM1 fiber</b> 150 m at 2 Gb/s		
9 micron single-mode fiber optic cable and long-wave SFPs	10 km ISL at 2 Gb/s		
9 micron single-mode fiber optic cable and extended-reach SFPs	35 km ISL at 2 Gb/s (B-series and M-series)		
Total distance <sup>2</sup>	200 km (B-series, C-series) 105 km (M-series)		
Fibre Channel using WDM <sup>3</sup>	100 km at 2 Gb/s 240 km at 1 Gb/s		

<sup>1</sup> Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

<sup>2</sup> Total distance is the sum of all cable segments.

<sup>3</sup> WDM distance is the maximum distance for the WDM link.

Table 82 (page 153) describes the distance rules for 1 Gb/s Fibre Channel connections when using 1 Gb/s Fibre Channel switch models.



**Table 82 Fibre Channel distance rules for 1 Gb/s switch models (B-series, C-series, and M-series switches)**

Interface/transport	Supported distances	Supported storage products
50 micron multi-mode fiber optic cable and short-wave GBICs and GLMs	500 m at 1 Gb/s	Heterogeneous SAN servers, Fibre Channel switches, and storage systems HP P6000 Continuous Access and HP P9000 (XP) Continuous Access EBS
62.5 micron multi-mode fiber optic cable and short-wave GBICs <sup>1</sup>	200 m at 1 Gb/s	
9 micron single-mode fiber optic cable and long-wave GBICs	10 km ISL at 1 Gb/s 35 km ISL (C-series, M-series)	
9 micron single-mode fiber optic cable and very long distance GBICs	100 km ISL (B-series) at 1 Gb/s	
Total distance <sup>2</sup>	200 km (B-series, C-series) 105 km (M-series)	
Fibre Channel using WDM <sup>3, 4</sup>	240 km at 1 Gb/s	

<sup>1</sup> Information for 62.5 micron fiber optic cable is provided to facilitate use of installed cable. HP recommends 50 micron fiber optic cable for new installations that require multi-mode fiber.

<sup>2</sup> Total distance is the sum of all cable segments.

<sup>3</sup> WDM distance is the maximum distance for the WDM link.

<sup>4</sup> Up to 240 km at 1 Gb/s, over a WDM link. Certain Fibre Channel switches may have reduced performance at this distance.

Table 83 (page 153) describes the distance rules for ATM extension Fibre Channel connections.

**Table 83 ATM extension Fibre Channel distance rules**

Interface/transport	Heterogeneous SAN host-to-disk storage systems	HP P9000 (XP) Continuous Access <sup>1, 2</sup>	EBS
One T1/E1 WAN per fabric	Not supported	Inter-site backbone with FC-to-IP converter, 100 ms IP network delay one-way or 200 ms round-trip	Not supported
One T1/E1 WAN per fabric (inverse multiplexing)		Not supported as one ISL in a single-fabric implementation	
T3/E3 WAN		Inter-site backbone with FC-to-IP converter, 100 ms IP network delay one-way or 200 ms round-trip	
Fractional and/or shared T3/E3 and OC 3 WAN			

<sup>1</sup> For supported IP bandwidth levels, see the product documentation.

<sup>2</sup> For more information on HP Continuous Access support limits, see “SAN extension” (page 249).

Table 84 (page 153) describes the distance rules for FCIP extension Fibre Channel connections.

**Table 84 FCIP extension IP network distance rules**

Interface/transport	HP P6000 Continuous Access and HP P9000 (XP) Continuous Access <sup>1</sup>	EBS <sup>2</sup>
HP B-series 400 MP Router, MP Router Blade (FCIP)	100 ms IP network delay one-way or 200 ms round-trip	100 ms IP network delay one-way or 200 ms round-trip
HP C-series MDS 9216i, IPS-8, 14/2, 18/4		
HP IP Distance Gateway	P6000: 100 ms IP network delay one-way or 200 ms round-trip XP: 50 ms IP network delay one-way or 100 ms round-trip	Not supported
HP MPX200 Multifunction Router FCIP	P6000: 100 ms IP network delay one-way or 200 ms round-trip	Not supported

**Table 84 FCIP extension IP network distance rules** (continued)

Interface/transport	HP P6000 Continuous Access and HP P9000 (XP) Continuous Access <sup>1</sup>	EBS <sup>2</sup>
	XP: Not supported	
<b>Legacy FCIP products</b>		
IPS-4	100 ms IP network delay one-way or 200 ms round-trip	Not supported
HP B-series MP Router (FCIP)		100 ms IP network delay one-way or 200 ms round-trip
HP SR2122-2		Not supported
Third-party SAN extension devices	See “SAN extension” (page 249).	Not supported

<sup>1</sup> For more information on HP Continuous Access support limits, see “SAN extension” (page 249).

<sup>2</sup> EBS is not supported with FCIP IPsec data encryption.

Table 85 (page 154) describes the distance rules for iSCSI bridging Fibre Channel connections.

**Table 85 iSCSI bridging Fibre Channel distance rules**

Interface/transport	Heterogeneous SAN Fibre Channel distances
HP B-series iSCSI Director Blade	See Table 78 (page 150) through Table 82 (page 153).
HP C-series MDS 9216i, IPS-4, IPS-8, 14/2, 18/4	
EVA iSCSI Connectivity Option	
HP MPX200 Multifunction Router iSCSI	
HP SR2122-2	

Table 86 (page 154) describes the distance rules for Fibre Channel routing connections.

**Table 86 Fibre Channel and IP network routing distance rules**

Interface/transport	Heterogeneous SAN Fibre Channel distances	HP P6000 Continuous Access and HP P9000 (XP) Continuous Access <sup>1</sup>	EBS
HP B-series 400 MP Router, MP Router Blade, and MP Router	See Table 80 (page 151) through Table 82 (page 153).	See Table 80 (page 151) through Table 82 (page 153) and Table 84 (page 153).	
HP C-series IVR (inter-VSAN routing)			

<sup>1</sup> For more information on HP Continuous Access support limits, see “SAN extension” (page 249).

## SAN fabric switch interoperability rules

HP supports three heterogeneous switch SAN configurations:

- Interoperable routed fabrics
- Dual interoperable, heterogeneous SAN fabrics
- Interoperable, heterogeneous switch fabrics

## Interoperable routed fabric

Connecting B-series and M-series edge fabrics through an MP Router allows for an interoperable routed fabric. You can share devices between these fabrics with the LSAN feature of the MP Router. For supported fabric modes and firmware versions, see “MP Router fabric rules” (page 107).

## Dual interoperable, heterogeneous SAN fabrics

A dual interoperable, heterogeneous SAN consists of two fabrics, with each fabric comprised exclusively of switches from a different series.

HP supports dual interoperable SAN fabrics for the following switch-series combinations:

- A SAN with one fabric containing B-series switches only and a second fabric containing C-series switches only
- A SAN with one fabric containing B-series switches only and a second fabric containing H-series switches only
- A SAN with one fabric containing B-series switches only and a second fabric containing M-series switches only

When creating a dual interoperable, heterogeneous SAN, consider the following:

- HP recommends that you use the same fabric topology and configuration in both fabrics to maintain balanced SAN performance.
- Design both fabrics using the lowest common denominator. For example, B-series switches support seven switch hops and M-series switches support three switch hops. Therefore, design both fabrics with a maximum of three switch hops.
- HA configurations require support for common HBA, driver, multipathing software, and storage array firmware versions because servers and storage connect to both fabrics.

---

**NOTE:** HP does not support MSA1000 with the MSA SAN Switch 2/8 in dual interoperable, heterogeneous SAN configurations.

---

## Interoperable, heterogeneous switch fabrics

An interoperable, heterogeneous switch fabric can contain different series of switches.

Table 87 (page 155) lists the switch combinations.

**Table 87 Heterogeneous switches in the same fabric**

Heterogeneous switch combinations	Reference
C-series and B-series	<i>Fabric Interoperability: Merging Fabrics Based on C-series and B-series Fibre Channel Switches Application Notes</i>
C-series and M-series	<i>Fabric Interoperability: Merging Fabrics Based on C-series and M-series Fibre Channel Switches Application Notes</i>
M-series and B-series	<i>Fabric Interoperability: Merging Fabrics Based on M-series and B-series Fibre Channel Switches Application Notes</i>

The documents referenced in Table 87 (page 155) are available on the HP SAN Infrastructure website:

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

---

**NOTE:** HP does not support MSA1000 storage systems or MSA1500 storage systems in interoperable, heterogeneous switch fabric configurations.

---

## Third-party switch support

HP Services offers support for certain third-party switches if you purchase third-party support through Multivendor Environment Services. For more information, see the HP Storage Services website:

[http://www8.hp.com/us/en/services/services-detail.html?pageTitle=Integrated-Multivendor-Services&compURI=tcm%3A245-807609&jumpid=reg\\_R1002\\_USEN](http://www8.hp.com/us/en/services/services-detail.html?pageTitle=Integrated-Multivendor-Services&compURI=tcm%3A245-807609&jumpid=reg_R1002_USEN)

HP Services also offers support for multivendor SANs. HP provides integrated hardware and software support with proactive problem prevention and 24x7 assistance.

## SAN performance considerations

The following SAN components affect SAN application performance:

- Host CPUs
- Fibre Channel HBAs
- SAN topology and the number of fabrics
- I/O transfer sizes and usage patterns
- RAID controllers
- Disk configuration

## Infrastructure factors

A single-switch fabric provides the highest level of performance. In a fabric with multiple switches, the following factors can affect performance:

- **Latency**

Switch latency is less than 5% (at 1 Gb/s) of the data transfer time; therefore, the number of switches and hops between devices is not a major performance factor. However, as devices send frames through more switches and hops, other data traffic in the fabric routed through the same ISL or path can cause oversubscription.

- **Oversubscription**

Oversubscription degrades Fibre Channel performance. When devices must contend for the same ISL or path, each device receives an equal share or  $1/n$ th of the available bandwidth on the path (where  $n$  is the number of contending devices). Oversubscription occurs when one or more devices sends more data than the total bandwidth available on the ISL or path.

- **Fabric interconnect speeds**

Fibre Channel supports 16 Gb/s, 8 Gb/s, 4 Gb/s, 2 Gb/s, and 1 Gb/s speeds. For optimum performance, configure a fabric with all components at the same, highest available speed. Additional factors such as distance, number of switch and device port buffers, and device response times can also affect performance.

- **Mixed Fibre Channel speeds**

For fabrics consisting of 16 Gb/s, 8 Gb/s, and 4 Gb/s; or 8 Gb/s, 4 Gb/s, and 2 Gb/s; or 4 Gb/s, 2 Gb/s, and 1 Gb/s switches and devices, the fabric segment connections negotiate the speed at which specific devices communicate.

The presence of lower speed devices in a fabric does not force other independent higher speed devices or paths to a lower speed. Fibre Channel requires that all 16 Gb/s ports be able to negotiate to 8 Gb/s and 4 Gb/s, all 8 Gb/s ports to 4 Gb/s and 2 Gb/s, and all 4 Gb/s ports to 2 Gb/s and 1 Gb/s speeds. Switch ports or user ports in a fabric communicate at the highest mutually supported speed.

## Performance guidelines

Although the topology and size of the fabric affect performance, adhering to the rules and recommendations outlined in this guide minimizes these factors. The topology designs have been defined to accommodate specific data access types. Recommendations on the number of ISLs based on device-to-device access ratios ensure that adequate bandwidth is available across the fabric, minimizing oversubscription.

To maximize fabric performance, HP recommends the following guidelines:

- Implement dual-fabric SANs.
- In a cascaded or core-edge fabric, position switches with the highest port speeds near the center of the fabric.
- Use the highest speed available for all infrastructure components and devices.
- Ensure that communicating devices have the same speed connectivity path through the fabric.
- Connect devices that communicate frequently to the same Fibre Channel switch.
- When possible, ensure that there is an equal number of high-bandwidth application servers and storage systems (for one-to-one access).
- Ensure that FCC is enabled on all C-series switches.

FCC allows C-series switches to intelligently regulate traffic across ISLs and ensure that each initiator-target pair of devices has the required bandwidth for data transfer. C-series switches can also prioritize frames using the QoS feature.

---

## Part III Host and storage system rules

Host and storage system rules are presented in these chapters:

- “Heterogeneous server rules” (page 159)
- “P2000 and MSA storage system rules” (page 203)
- “P6000/EVA storage system rules” (page 213)
- “P9000/XP and VA storage system rules” (page 227)
- “SVSP storage system rules” (page 234)
- “HP 3PAR storage system rules” (page 241)
- “Enterprise Backup Solution” (page 247)

---

## 10 Heterogeneous server rules

This chapter describes platform configuration rules for SANs with specific operating systems and heterogeneous server platforms:

- “SAN platform rules” (page 160)
- “Heterogeneous storage system support” (page 160)
- “Blade servers” (page 161)
- “HP c-Class BladeSystem Virtual Connect” (page 161)
- “BladeSystem with Brocade Access Gateway mode” (page 163)
- “BladeSystem with Cisco N\_Port Virtualization mode” (page 165)
- “NPV with FlexAttach” (page 168)
- “HP BladeSystem c3000 enclosure considerations” (page 169)
- “HBA N\_PortID Virtualization” (page 170)
- “NonStop servers (XP only)” (page 171)
- “HP-UX SAN rules” (page 183)
- “HP OpenVMS SAN rules” (page 185)
- “HP Tru64 UNIX SAN rules” (page 187)
- “Apple Mac OS X SAN rules” (page 188)
- “IBM AIX SAN rules” (page 190)
- “Linux SAN rules” (page 191)
- “Microsoft Windows SAN rules” (page 193)
- “Sun Solaris SAN rules” (page 196)
- “VMware ESX SAN rules” (page 198)
- “Citrix Xen SAN rules” (page 199)
- “Heterogeneous SAN storage system coexistence” (page 200)
- “Server zoning rules” (page 202)

The platform configuration rules in this chapter apply to SANs that comply with the following fabric guidelines:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “H-series switches and fabric rules” (page 131)
- “M-series switches and fabric rules” (page 140)

Before implementation, contact an HP storage representative for support information for specific configurations, including the following elements:

- Server model
- Storage system firmware
- SAN attachment
- HBAs and drivers
- Multipathing

# SAN platform rules

Table 88 (page 160) describes SAN platform rules for all SAN server configurations.

**Table 88 General SAN platform rules**

Rule number	SAN platform configuration
1	Any combination of heterogeneous clustered or standalone servers with any combination of storage systems is supported. The configuration must conform to requirements and rules for each SAN component, including: <ul style="list-style-type: none"><li>• Operating system</li><li>• Fabric</li><li>• Storage system</li><li>• Mixed storage system types</li></ul>
2	All HP and multivendor hardware platforms and operating systems that are supported in a homogeneous SAN are also supported in a heterogeneous SAN.  In a heterogeneous SAN, define zones by operating system. Storage systems can be in multiple operating system type zones.
3	Servers can connect to multiple fabrics. The number of supported fabrics per server depends on the maximum number of Fibre Channel HBAs supported for the server. See <a href="#">“EVA single-server maximum configurations”</a> (page 217).

For cabling options for platforms that support high-availability multipathing, see [“Cabling”](#) (page 222).

## Heterogeneous storage system support

HP supports HP storage products on shared hosts and HBAs in HP fabric environments that also have third-party storage products. A third-party cooperative support agreement between HP Services and the third party is required if HP will provide a single support point of contact that includes the third-party storage. HP provides technical support for its products and cooperates with the third party's technical support staff, as needed.

HP provides best-practices recommendations for connecting devices in the SAN. See [“Best practices”](#) (page 372). These rules apply to configurations that include HP SAN storage products and heterogeneous third-party SAN storage products:

- Use zones to isolate HP storage ports from third-party storage ports.
- HP storage zones are governed by HP product-specific configuration guidelines. See the following HP storage system chapters:
  - [“P2000 and MSA storage system rules”](#) (page 203)
  - [“P6000/EVA storage system rules”](#) (page 213)
  - [“P9000/XP and VA storage system rules”](#) (page 227)
  - [“SVSP storage system rules”](#) (page 234)
  - [“HP 3PAR storage system rules”](#) (page 241)
- Overlapping zones with multiple multi-vendor storage ports are not supported.
- Third-party storage zones are governed by product-specific configuration guidelines (see the third-party product documentation).

For third-party fabric and switch support, see [“Third-party switch support”](#) (page 156).

For storage system coexistence support, see [“Heterogeneous SAN storage system coexistence”](#) (page 200).



## Blade servers

Table 89 (page 161) lists supported blade servers and switches.

**Table 89 Supported blade servers**

Blade servers	Switches
BL460c BL465c BL480c BL685c BL860c	B-series Brocade 4Gb SAN Switch for HP c-Class BladeSystem C-series Cisco MDS 9124e Fabric switch for HP c-Class BladeSystem H-series M-series
BL2xp BL45p BL60p	B-series Brocade 4Gb SAN Switch for HP p-Class BladeSystem C-series M-series McDATA 4Gb SAN Switch for HP p-Class BladeSystem
BL3xp	B-series Brocade 4Gb SAN Switch for HP p-Class BladeSystem C-series McDATA 4Gb SAN Switch for HP p-Class BladeSystem
BL40p	B-series C-series M-series

For HP ProLiant BL c-Class SAN product support, see the support matrix at <http://h18006.www1.hp.com/storage/saninfrastructure/switches/b4gbsscblade>.

For HP ProLiant BL p-Class SAN product support, see <http://h18006.www1.hp.com/products/storageworks/bladessystemmatrix/index.html>.

For fabric rules when using the HP p-Class BladeSystem SAN switches, see:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “M-series switches and fabric rules” (page 140)
- “SAN fabric connectivity and switch interoperability rules” (page 146)

For third-party fabrics, HP supports the blade server, the HP p-Class BladeSystem SAN switch, and switch connectivity (including the switch interface) to the non-HP fabric. See “Third-party switch support” (page 156).

## HP c-Class BladeSystem Virtual Connect

The HP 4Gb VC-FC module is an optional Fibre Channel interconnect offered on HP BladeSystem c-Class servers. VC-FC employs NPIV protocol to aggregate multiple Fibre Channel HBAs over a smaller number of N\_Port uplinks. The VC-FC module is Fibre Channel standards based and is compatible with all NPIV standards-compliant switch products.

The VC-FC module can aggregate up to sixteen Fibre Channel HBA ports at the back end connected through the enclosure backplane, and comes with four uplinks or N\_Ports which connect to an external Fibre Channel fabric. Each of the four uplinks is configurable to aggregate between zero and sixteen Fibre Channel HBAs resulting in an over-subscription ratio of up to 16:1. The default

over-subscription ratio is 4:1 with other user configurable options of 8:1 and 16:1. The VC-FC module appears as a pass-thru device to the network. Any changes to the server are transparent to its associated network. This separates the servers from the fabric and relieves SAN administrators from server maintenance.

The VC-FC module is one of two components of the HP Virtual Connect technology. The other component of the solution is the HP Virtual Connect Ethernet (VC-Enet) module. Both modules are managed by Virtual Connect manager software embedded on VC-Enet module. The HP Virtual Connect solution requires at least one VC-Enet module and one VC-FC module.

## HP Virtual Connect advantages

Advantages of HP Virtual Connect include:

- Simplifies server connections by separating server management from LAN and SAN management. SANs are managed at the core switch, separate from blade enclosure.
- Allows you to change servers in minutes, not days. You can add, replace, and recover servers quickly without affecting SANs or LANs.
- Provides connectivity to multiple fabric types (B-series, C-series, H-series, M-series), without interoperability constraints.
- Does not use a domain ID, preventing issues with the switch count scalability limit in a fabric and eliminating the additional hop between switches.
- Reduces the number of cables and SFPs compared to a Pass-thru solution.

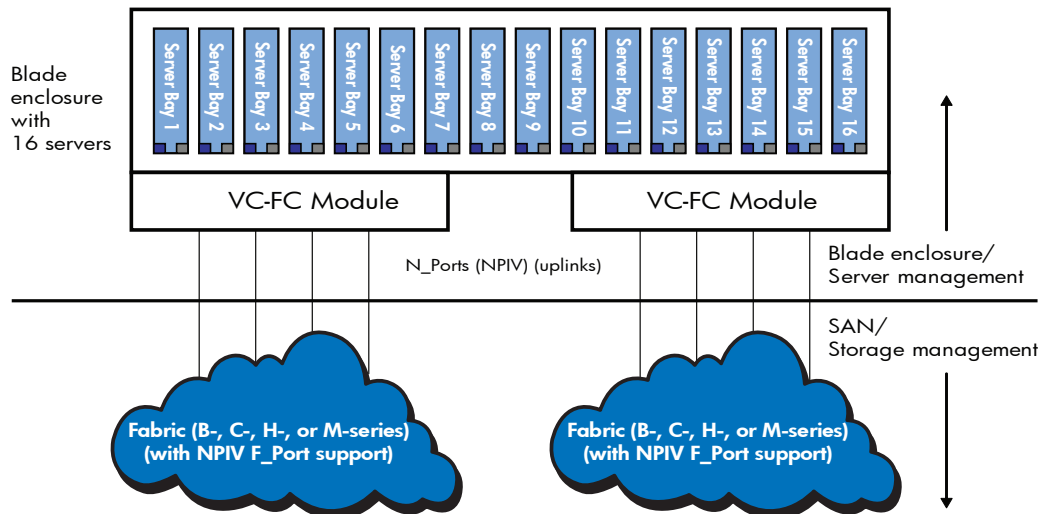
## HP VC-FC connectivity guidelines

Deploy HP VC-FC in environments where you need to manage servers without impacting SAN management (that is, the server administrator manages the entire configuration). There are several customer configurations with varying numbers of VC-FC modules and blade enclosures. The actual configurations depend on customer connectivity requirements, availability of existing equipment, and future growth requirements. [Table 90 \(page 162\)](#) describes the HP VC-FC rules and [Figure 54 \(page 163\)](#) shows an HP VC-FC example.

**Table 90 HP Virtual Connect Fibre Channel guidelines and rules**

Rule number	Description
1	For high availability, configure two redundant fabrics and two VC-FC modules, with each HBA connecting to one fabric through one VC-FC module.
2	VC-FC is supported with B-series, C-series, H-series, and M-series fabrics. The VC-FC module must connect to a switch model that supports NPIV F_Port connectivity. Certain switch models may require a license.
3	HP supports a maximum of 4 VC-FC modules per blade enclosure.

Figure 54 HP Virtual Connect Fibre Channel configuration



## BladeSystem with Brocade Access Gateway mode

AG mode is a software-enabled feature available with the Brocade 8Gb SAN Switch for HP BladeSystem c-Class and the Brocade 4Gb SAN Switch for HP c-Class and p-Class BladeSystems with Fabric OS 5.2.1b (or later). AG mode does not require the purchase of additional hardware or software.

Blade switches in AG mode function as port aggregators using NPIV to connect to NPIV-compliant Fibre Channel switches (including other vendor switches). The blade switches are logically transparent to the hosts and fabric—they no longer function as standard switches.

**NOTE:** This section discusses HP c-Class BladeSystems, but HP p-Class BladeSystems also support AG mode.

The Brocade 8Gb SAN Switch for HP BladeSystem c-Class and the Brocade 4Gb SAN Switch for HP c-Class BladeSystems in AG mode supports a maximum of 24 ports:

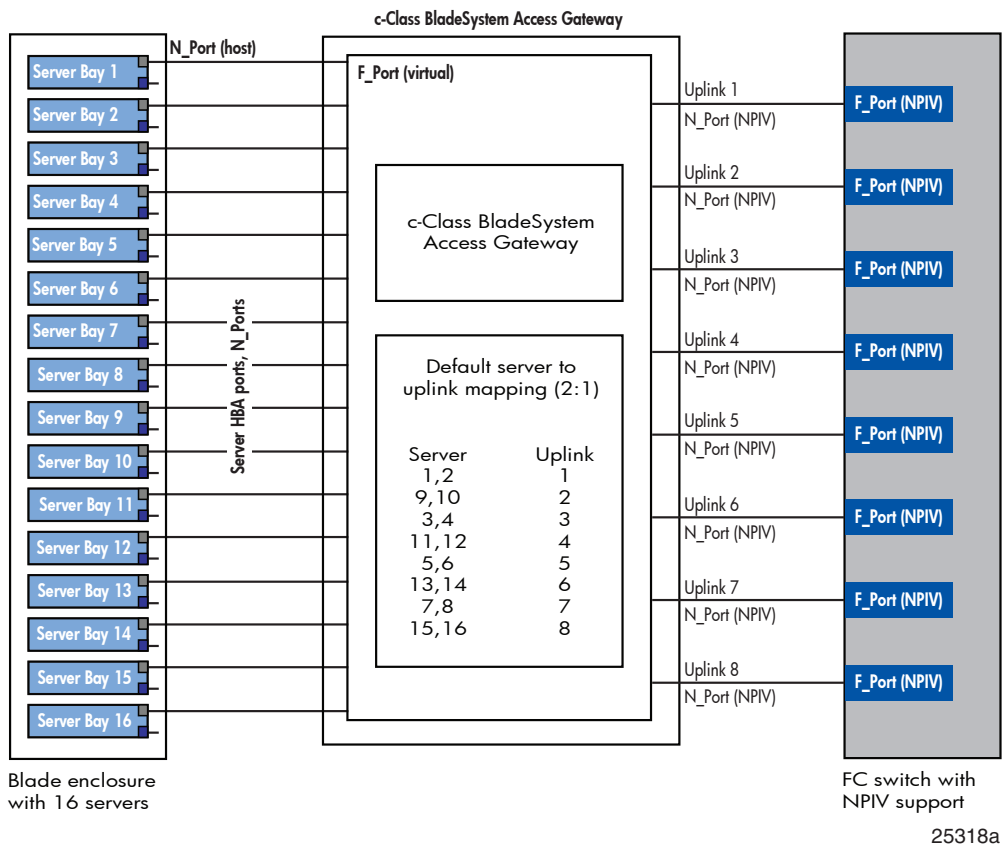
- Maximum of 16 ports for back-end connections to blade server
- Maximum of 8 external ports used as uplink N\_Ports

AG mode features include:

- The 8 external ports function as N\_Ports, supporting NPIV. They connect to standard switches that support NPIV-compliant F\_Ports.
- AG mode does not use a domain ID, preventing domain-count limits in large fabrics.
- AG mode uses port mapping between the host-facing ports (virtual F\_Ports) and the external uplink ports (N\_Ports). The default mapping is 2:1 ports, which you can reconfigure as needed.

Figure 55 (page 164) shows a view of an HP c-Class BladeSystem in AG mode.

**Figure 55 Brocade 4Gb SAN Switch for HP c-Class BladeSystem in Access Gateway mode**



**NOTE:** The uplink ports (N\_Ports) in Figure 55 (page 164) are from the AG, not the hosts.

## Failover policy and failback policy

AG mode supports the failover and failback policies (enabled by default), which you can configure on a per-port basis.

The failover policy enables automatic remapping of hosts to other online N\_Ports if N\_Ports go offline. It evenly distributes hosts among the available N\_Ports. This policy ensures a smooth transition with minimal traffic disruption when a link fails between an N\_Port on the AG and an F\_Port on the external fabric.

The failback policy automatically routes hosts back to the original N\_Ports when they come online.

**NOTE:** When a failover or failback occurs, hosts must be logged back in to resume I/O.

## AG mode considerations

AG mode considerations follow:

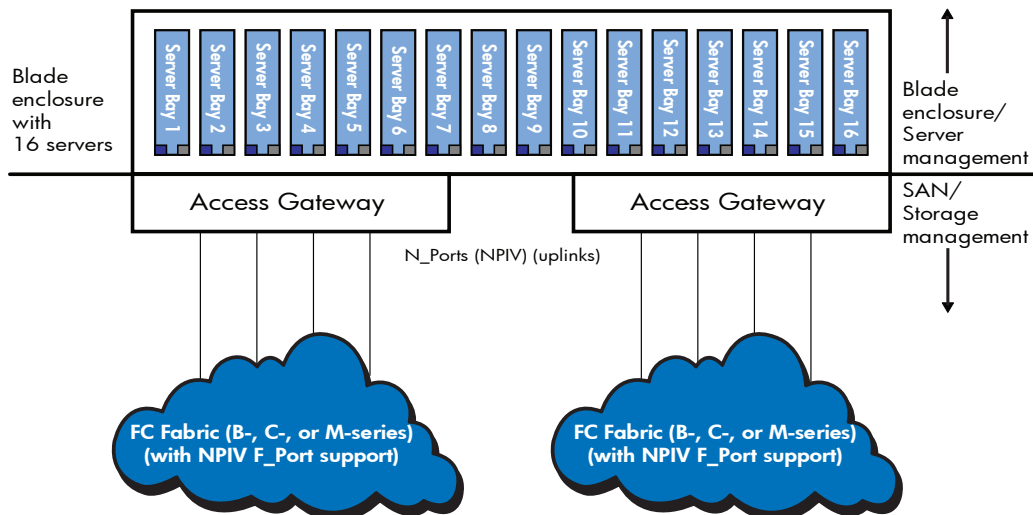
- Ability to connect B-, C-, and M-series fabrics without interoperability constraints (for support information, see the release notes)
- Flexible licensing option (12 or 24 ports, with a 12-port upgrade option on the 12-port model)
- Ability to use in either switch mode or AG mode (cannot function in both modes simultaneously, software selectable)
- Share Fabric OS with B-series switches
- Port failover between N\_Ports (uplinks)
- Reduces the number of cables and SFPs compared to a Pass-thru solution

- No SAN management from the BladeSystem enclosure once the initial connections have been configured
- No direct storage attachment (requires at least one external Fibre Channel switch)
- Lacks Fibre Channel embedded switch features (ISL Trunking, dynamic path selection, and extended distances) with external links from AG to core switches
- Managed separately from the BladeSystem, but if used with B-series switches, uses common Fabric OS
- Cannot move servers without impacting the SAN (Virtual Connect feature not available)

## AG mode connectivity guidelines

AG-based solutions are best suited for B-series-only fabrics where you want multivendor switch interoperability through N\_Ports instead of E\_Ports. [Figure 56 \(page 165\)](#) shows an AG with dual redundant fabrics.

**Figure 56 Access Gateway with dual redundant fabrics**



25317b

**NOTE:** The N\_Ports in [Figure 56 \(page 165\)](#) are not host N\_Ports and cannot be connected directly to storage.

Configuration highlights for [Figure 56 \(page 165\)](#) include:

- Redundant SANs, with each server connecting to one fabric through one AG module
- Ability to connect to B-series, C-series, and M-series fabrics
- Support for up to six AGs per blade enclosure

## BladeSystem with Cisco N\_Port Virtualization mode

NPV mode is a software-enabled feature available on the Cisco MDS 9124e Fabric Switch for HP c-Class BladeSystem with SAN-OS 3.2(1a) (or later). NPV mode does not require the purchase of additional hardware or software.

NPV is available only if the Cisco MDS 9124e is in NPV mode; if the fabric switch is in switch mode, NPV is not supported. To use NPV, the end devices connected to a switch in NPV mode must log in as N\_Ports. All links from the end switches in NPV mode to the core switches are established as NP\_Ports, not E\_Ports for ISLs.

An NP\_Port is an NPIV uplink from the NPV device to the core switch. Switches in NPV mode use NPIV to log in multiple end devices that share a link to the core switch. The Cisco MDS 9124e Fabric Switch is transparent to the hosts and fabric—they no longer function as standard switches.

**NOTE:** This section describes HP c-Class BladeSystems. NPV mode is also supported on the Cisco MDS 9124 and MDS 9134 Fabric Switches. For more information, see the *Cisco MDS 9000 Configuration Guide*.

The Cisco MDS 9124e Fabric Switch in NPV mode supports a maximum of 24 ports:

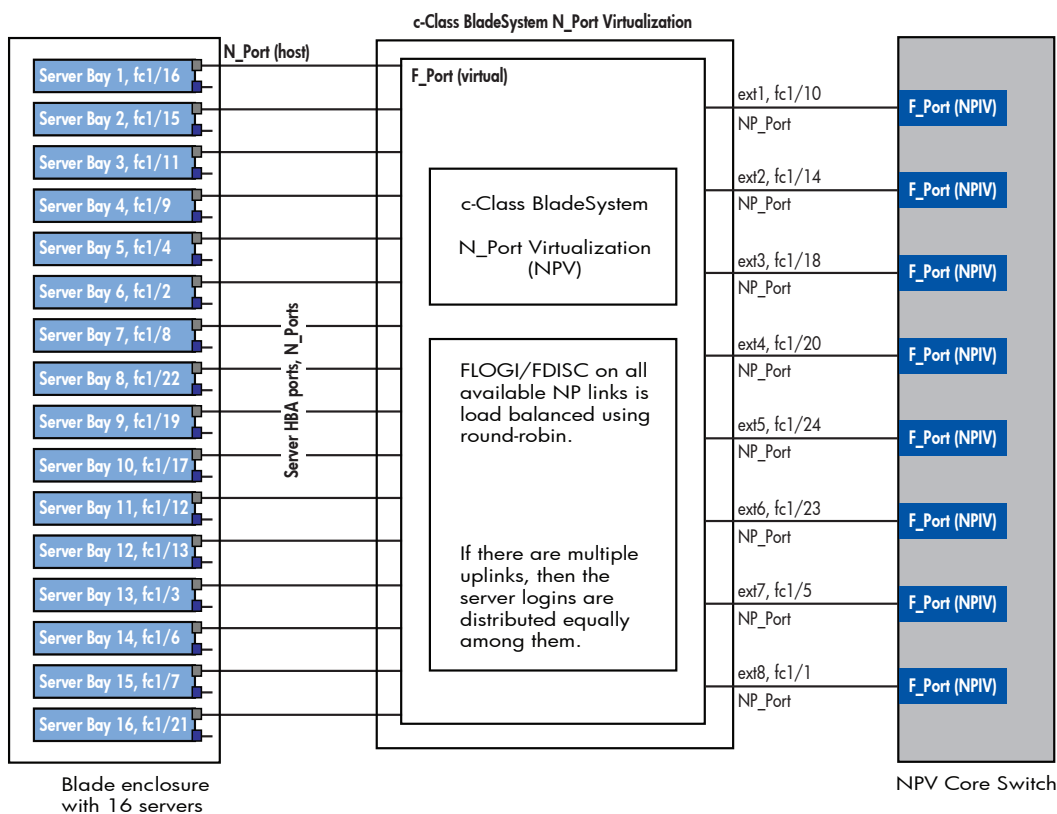
- Up to 16 ports for back-end connections to the BladeSystems
- Up to 8 external ports used as uplink ports (NP\_Ports)

NPV mode features include:

- Eight external ports used as NP\_Ports, supporting NPIV. These ports connect to standard switches (including vendor switches) that support NPIV-compliant F\_Ports.
- Does not use domain IDs, removing any domain-count limitations in large fabrics.
- Port mapping between the host-facing ports (virtual F\_Ports) and the external uplink ports (NP\_Ports).

Figure 57 (page 166) shows an HP c-Class BladeSystem in NPV mode.

**Figure 57 Cisco MDS 9124e Fabric Switch for HP c-Class BladeSystem in NPV mode**



25345a

**NOTE:** The NP\_Ports in Figure 57 (page 166) are on the NPV devices, not the hosts.

## Failover policy

The failover policy enables automatic remapping of hosts if NP\_Ports go offline. It evenly distributes the hosts among the available NP\_Ports. This policy ensures a smooth transition with minimal traffic

disruption when a link fails between an NP\_Port on the NPV devices and an F\_Port on the external fabric. To avoid disruption when an NP\_Port goes online, the logins are not redistributed.

## NPV mode considerations

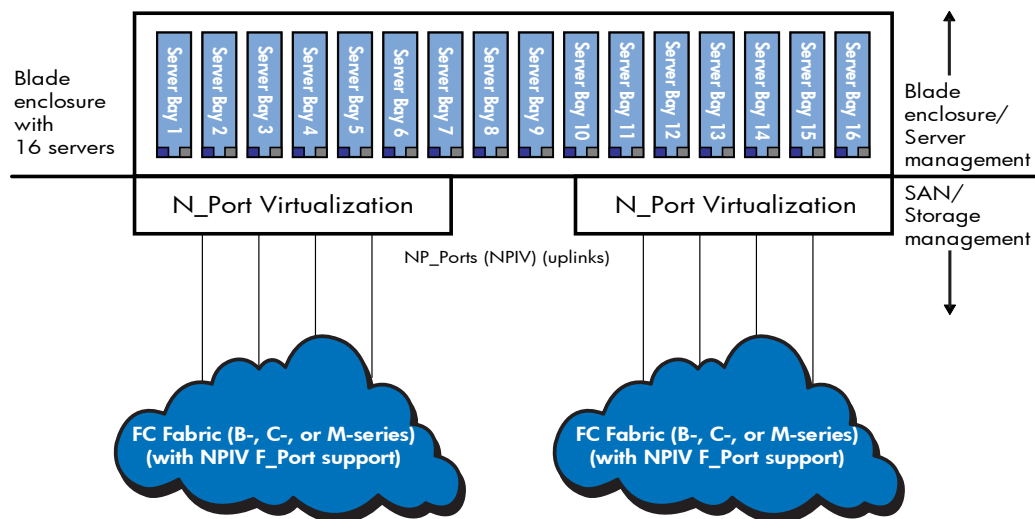
Consider the following:

- Nondisruptive upgrades are supported.
- Grouping devices into different VSANs is supported.
- A load-balancing algorithm automatically assigns end devices in a VSAN to one of the NPV core switch links (in the same VSAN) at initial login.
- You can connect B-, C-, and M-series fabrics without interoperability constraints (for support information, see the release notes).
- A flexible licensing option is available (12 or 24 ports, with a 12-port upgrade option on the 12-port model).
- You can select to use either switch mode or NPV mode.
- Failover between NP\_Ports (uplinks) is supported.
- Direct storage attachment is not supported (requires at least one external Fibre Channel switch).
- F\_Ports, NP\_Ports, and SD\_Ports are supported.
- NPIV-capable module servers (nested NPIV) are supported.
- Local switching is not supported. All traffic is switched using the NPV core switch.
- Remote SPAN is not supported.
- NPV mode is managed separately from the BladeSystem; however, if used with C-series switches, it uses a common SAN-OS.

## NPV mode connectivity guidelines

NPV solutions are best suited for C-series-only fabrics in which you want multivendor switch interoperability through NP\_Ports instead of E\_Ports. [Figure 58 \(page 167\)](#) shows an NPV device with dual redundant fabrics.

**Figure 58 NPV device with dual redundant fabrics**



25346b

---

**NOTE:** The NP\_Ports in [Figure 58 \(page 167\)](#) are not host N\_Ports and cannot connect directly to storage.

---

The configuration shown in [Figure 58 \(page 167\)](#) includes:

- Redundant SANs, with each server connecting to one fabric through one NPV device
- Connectivity to C-series, B-series, and M-series fabrics
- Support for up to six NPV devices per HP BladeSystem c7000 enclosure, or three NPV devices per HP BladeSystem c3000 enclosure

## NPV with FlexAttach

The Cisco MDS 9124e Fabric Switch for HP c-Class BladeSystem, MDS 9124 switch, and MDS 9134 switch support NPV with FlexAttach. FlexAttach provides automatic mapping of physical WWNs to virtual WWNs using NAT. When NPV mode is enabled, FlexAttach allows SAN and server administrators to install and replace servers without having to rezone or reconfigure the SAN.

With FlexAttach, you can perform the following tasks without the need to make SAN or storage configuration changes:

- **Preconfiguration**—You can preconfigure the SAN for the addition of new servers whose WWPNs are unknown, using the virtual WWPNs. Once the servers are available, you can bring them online and into the fabric.
- **Replacement (new server)**—You can replace an existing server with a new server. FlexAttach assigns a virtual WWPN to the server port.
- **Replacement (spare server)**—You can bring a spare server online by moving the virtual WWPN from the current server port to the spare server port.
- **Server redeployment**—You can move a server to a different NPV switch (in the same fabric or VSAN). FlexAttach allows you to manually create and transfer virtual WWPNs from one server port to another server port.

---


**NOTE:** Other tasks may require configuration changes. For more information about FlexAttach, see the *Cisco MDS 9000 Family CLI Configuration Guide*. The terms pWWN and WWPN are used interchangeably.

---

[Figure 59 \(page 169\)](#) shows a view of an HP c-Class BladeSystem using NPV with FlexAttach.



**Figure 59 Cisco MDS 9124e Fabric Switch for HP c-Class BladeSystem using NPV with FlexAttach**



**NOTE:** The names of the uplink ports (N\_Ports 1 through 8) in Figure 59 (page 169) are symbolic only. See the NPV documentation for the actual port numbers.

## HP BladeSystem c3000 enclosure considerations

Consider the following when using the HP BladeSystem c3000 enclosure:

- The c3000 has four interconnect bays: 1, 2, 3, and 4. If Fibre Channel switch redundancy is required, use interconnect bays 3 and 4.
- Interconnect bay 1 is dedicated to Ethernet or NIC connections; it cannot be used for Fibre Channel connections.
- Interconnect bay 2 can be used for Ethernet, NIC, or Fibre Channel connections; it is accessible through the mezzanine 1 card only. If you use Fibre Channel connections from mezzanine 1 cards, connect them to the interconnect bay 2 switch only, which provides port redundancy but not switch redundancy.
- Interconnect bay 2 cannot be used for VC-FC; it is restricted to Ethernet, NIC, or Fibre Channel connections. VC-FC modules must use interconnect bays 3 and 4.
- Interconnect bays 3 and 4 can be used for Fibre Channel connections and switch redundancy. The full-height or half-height mezzanine 2 cards provide Fibre Channel port and switch redundancy.

## HBA N\_PortID Virtualization

HBA NPIV is a Fibre Channel standard that allows multiple N\_Ports to connect to a switch F\_Port. HBA NPIV is used on servers running a VOS. You can assign a unique virtual port name to each VM that shares the HBA. NPIV is supported on all 8 Gb and 4 Gb Emulex and QLogic HBAs when using the vendor-supplied VOS drivers.

### HBA NPIV considerations

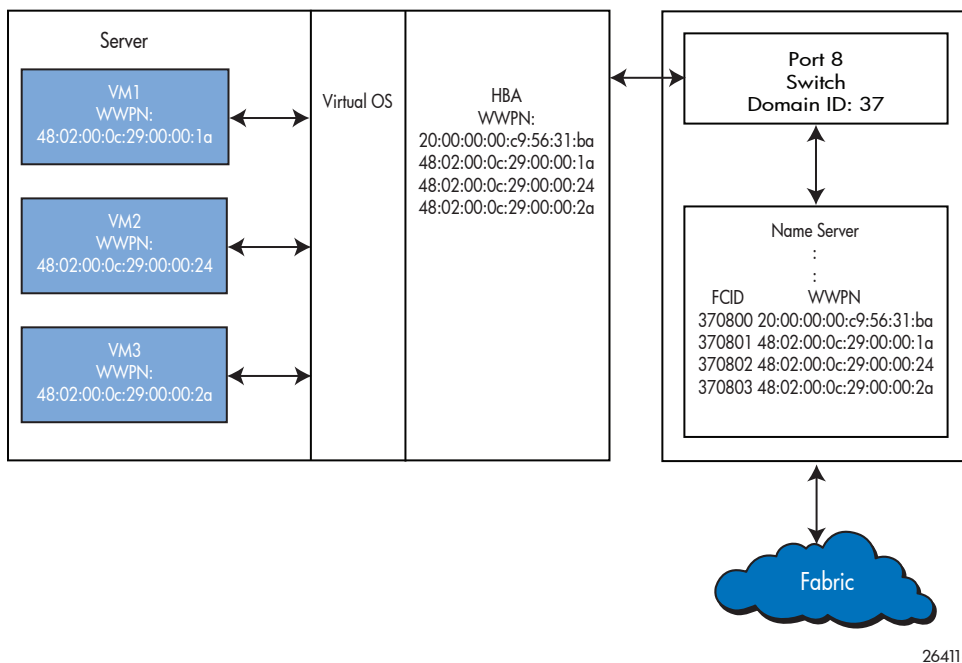
Consider the following points when implementing a SAN with VOS servers using HBA NPIV:

- You can assign and manage the virtual WWPN through the VOS. WWPN provides increased security and integrity because you can create discreet zones based on the port name.
- You must verify that the WWPNs in the SAN are unique. This is especially important for complex SANs with heterogeneous VOSs.
- You may need to enable HBA NPIV for some HBA and VOS combinations.
- F\_Port NPIV support differs for B-series, C-series, H-series, and M-series switches. For information about setting up switches for use with HBA NPIV, see the switch documentation.
- VMware ESX 3.5 and 4.0 are the only VOSes with native support for HBA NPIV. The supplied Emulex and QLogic drivers are NPIV enabled by default.
- Each VOS may have restrictions or requirements for HBA NPIV. For information about setting it up for use with HBA NPIV, see the operating system documentation.
- If a VOS supports VM migration, the virtual WWPNs associated with the VM will migrate.

### HBA NPIV connectivity guidelines

Figure 60 (page 171) shows the logical relationship between virtual WWPNs and a VOS with HBA NPIV enabled. A server running a VOS has three instances of VMs. The server has an HBA with a manufacturing-assigned WWPN (20:00:00:00:c9:56:31:ba), and is connected to port 8 of a switch whose domain ID is 37. The VOS generates three virtual WWPNs and maps them to the VMs. The VOS uses an operating system-specific algorithm to create the WWPNs, which can include a registered vendor unique ID.

**Figure 60 VOS with HBA NPIV enabled**



When using HBA NPIV, consider the following:

- When a VOS initializes, the HBA performs a fabric login using the manufacturing-assigned WWPN, and the switch assigns a FCID for the login session. The HBA WWPN and associated FCID are logged in the fabric name server.

When an NPIV-enabled VM initializes, the HBA performs another fabric login using the virtual WWPN associated with that VM, which creates another FCID and entry in the fabric name server. This process is repeated for each NPIV-enabled VM.

- When a VM stops, the entry is removed from the fabric name server. The relationship between FCIDs assigned to multiple N\_Ports logged in on the same F\_Port is not defined by the standards; instead, the switch vendors provide implementation details. In [Figure 60 \(page 171\)](#), the FCIDs have common values for the WWPN domain and area fields, and the port field value is incremented for each new login.

## NonStop servers (XP only)

NonStop servers are supported in direct host attach and SAN configurations for specific storage systems.

NonStop servers	Storage systems
<p>S-series servers:</p> <ul style="list-style-type: none"> <li>S760, S76000</li> <li>S78, S780, S7800, S78000</li> <li>S86000, S88000</li> </ul> <p>NS-series servers:</p> <ul style="list-style-type: none"> <li>NS1000, NS1200</li> <li>NS14000, NS14200</li> <li>NS16000, NS16200</li> </ul> <p>NonStop BladeSystem servers:</p> <ul style="list-style-type: none"> <li>NB50000c</li> </ul>	<p>XP disk arrays:</p> <ul style="list-style-type: none"> <li>XP128, XP1024 (RAID450)</li> <li>XP10000, XP12000 (RAID500)</li> <li>XP20000, XP24000 (RAID600)</li> </ul>

There are three types of I/O interfaces used to connect NonStop servers to XP disk arrays. See [Table 94 \(page 173\)](#).

- **FCSAs**—These are the NonStop version of Fibre Channel HBAs, used to connect to XP disk arrays. The FCSA module slides into an 11U, rack-mounted Input/Output Adapter Module Enclosure (IOAME), which can hold up to 10 ServerNet I/O Adapters, either FCSAs or Gigabit Ethernet 4-port ServerNet Adapters (G4SAs), for Ethernet connectivity.
- **VIO**—NS1000, NS1200, NS14000, and NS14200 servers use the VIO instead of the IOAME. The VIO interface consists of two 4U VIO enclosures per system, one each for the X and Y ServerNet fabrics. Each VIO enclosure has four embedded Fibre Channel ports in a Fibre Channel PIC, for a total of eight embedded Fibre Channel ports in the NonStop system. (The Fibre Channel ports provide the same functionality as the FCSAs in IOAME systems.) Each VIO enclosure can be expanded to eight Fibre Channel ports using optional PICs in the expansion slots, for a total of 16 Fibre Channel ports in the NonStop system. The expanded ports can be used for FCDM or HP Enterprise Storage System (ESS) connections.
- **Storage CLIM**—CLIMs are only available on the HP Integrity NonStop NB50000c BladeSystem. CLIMs are rack mounted in the NonStop server cabinet and connect to one or two X- and Y-fabric ports through fiber cables running from ServerNet PICs to ServerNet ports on the NonStop server. CLIMs provide the physical interface to storage devices and support SAS and FC connections. CLIMs also perform certain storage management tasks previously done on the NonStop server. For information about differences between configuring the storage subsystem on IOAME-based systems and on CLIM-based systems, see the *NonStop Cluster I/O Module (CLIM) Installation and Configuration Guide*.

**NOTE:** Consider the following VIO requirements:

- For NS1000 and NS1200 servers, expanded ports are available only to customers who have the HP ESS.
- The VIO enclosure software is not backward compatible and is supported only on H06.08 and later RVUs.
- Prior to December 2006, the NS1000 and NS14000 servers were shipped with a limited IOAME configuration known as the IO Core, which consisted of an IOAME with six adapter slots rather than the usual ten slots. Customer installations with the IO Core configuration will be supported until December 2011.

[Table 91 \(page 172\)](#) describes supported NonStop server configurations with IOAMEs.

**Table 91 NonStop high-availability configurations using IOAMEs**

NonStop server support	Direct host attach configuration minimum/recommended	SAN configuration minimum/recommended	Maximum availability configuration minimum/recommended
Number of Fibre Channel SAN fabrics	0/0	2/2	2/4
Number of XP storage systems	1/1	1/1	1/2
Number of IOAMEs	1/1	1/1	2/2
Number of Fibre Channel ServerNet adapters	2/4	2/4	4/4

Table 92 (page 173) describes supported NonStop server configurations with VIO enclosures.

**Table 92 NonStop high-availability configurations using VIO enclosures (NS1000, NS1200, NS14000, and NS14200 only)**

NonStop server support	Direct host attach configuration minimum/recommended	SAN configuration minimum/recommended	Maximum availability configuration minimum/recommended
Number of Fibre Channel SAN fabrics	0/0	2/2	2/4
Number of XP storage systems	1/1	1/1	1/2
Number of VIO enclosures per server	2/2	2/2	2/2
Number of 4-port Fibre Channel PICs per VIO enclosure	1/2	1/2	2/2

Table 93 (page 173) describes supported NonStop BladeSystem with XP configurations using CLIMs.

**Table 93 NonStop BladeSystem with XP high-availability configurations using CLIMs (HP Integrity NonStop NB5000c BladeSystem only)**

NonStop server support	Direct host attach configuration minimum/recommended	SAN configuration minimum/recommended	Maximum availability configuration minimum/recommended
Number of Fibre Channel SAN fabrics	0/0	2/2	2/4
Number of XP storage systems	1/1	1/1	1/2
Number of CLIMs per Integrity NonStop BladeSystem	2/4	2/4	2/4
Number of dual-port Fibre Channel HBAs per CLIM	1/2	1/2	2/2

Table 94 (page 173) describes supported I/O modules for XP connectivity with specific NonStop systems.

**Table 94 Supported I/O modules for XP connectivity with specific NonStop systems**

NonStop servers	CLIM enclosure	IOAME with FCSAs installed	VIO enclosure
NonStop BladeSystems	Yes	Yes	No
NS16000, NS16200	No	Yes	No
NS14000, NS14200	No	Before December 2006 NS14000 only	After December 2006 NS14000 and NS14200
NS1000, NS1200	No	Before December 2006 NS1000 only	After December 2006 NS1000 and NS1200

Table 95 (page 174) describes storage system configuration rules for NonStop servers.

**Table 95 NonStop server configuration rules**

Rule number	Description
1	Requires a minimum of one XP storage system for storage connectivity.
2	Requires a minimum of one IOAME on the server. For the NS1000, NS1200, NS14000, and NS14200 servers using VIO, two VIO enclosures are used instead of the IOAME. For BladeSystems using CLIMs, two CLIMs are used instead of the IOAME.
3	Requires a minimum of two FCSAs in an IOAME, as shown in <a href="#">Figure 61 (page 176)</a> , <a href="#">Figure 64 (page 177)</a> , <a href="#">Figure 67 (page 179)</a> , and <a href="#">Figure 70 (page 180)</a> . Each FCSA has two Fibre Channel ports.
4	Servers using VIO require one embedded 4-port Fibre Channel PIC per VIO enclosure (total of two PICs per server) for basic connectivity, as shown in <a href="#">Figure 62 (page 176)</a> and <a href="#">Figure 65 (page 178)</a> . Logical and physical redundancy of the storage system, as shown in <a href="#">Figure 68 (page 179)</a> and <a href="#">Figure 71 (page 181)</a> , requires the addition of one expansion 4-port Fibre Channel PIC per VIO enclosure (total of two expansion PICs per server). For optimal I/O performance, use a separate I/O path for each logical disk volumes.
5	The following restrictions apply to CLIMs: <ul style="list-style-type: none"> <li>• The maximum number of LUNs for each CLIM, (including SAS disks, XP disk arrays, and tapes) is 512. Each primary, backup, mirror, and mirror backup path is counted.</li> <li>• The XP LUN range for each port is 0 to 499.</li> <li>• The maximum number of XP ports for each CLIM is four.</li> <li>• The maximum number of mirrored XP volumes is 256 with two CLIMs, and 512 with four CLIMs.</li> </ul>
6	Each LDEV requires two LUNs.
7	No boot support.
8	Contact HP support for host mode requirements on storage system ports.
9	For more information on supported B-series and C-series switches, contact HP support.
10	Each fabric must contain switches of the same series only—switches from multiple series within a fabric is not supported.
11	Slot 1 of the VIO enclosure contains a 4-port Fibre Channel PIC, which provides FCSA functionality of the IOAME systems. You can double the number of Fibre Channel ports on a server (from 8 to 16) by adding a 4-port Fibre Channel PIC in slot 7c of each VIO enclosure. This requires the installation of an expansion board in slot 7b of each VIO enclosure. FCDM or ESS connections can make use of these expanded Fibre Channel ports. <b>Note:</b> For NS1000 and NS1200 servers, expanded ports are available only to customers who have the HP ESS.
12	The VIO enclosure software is not backward compatible. This product is supported only on H06.08 and later RVUs.
13	The CLIM software is not backward compatible. This product is supported only on J06.04 and later RVUs.
14	<b>Direct host attach</b> <ul style="list-style-type: none"> <li>• HP recommends host-based mirroring. For example, each LDEV (P) is mirrored to a separate LDEV (M) on separate XP ports (p, b, m, mb paths are used). A nonmirrored volume is allowed. For example, each LDEV (P) is not mirrored to a separate LDEV (M) on separate XP ports (only p and b paths are used).</li> <li>• For high availability, the primary (P) LDEVs and mirror (M) LDEVs must be configured on separate array ACP pairs.</li> <li>• For high availability, the p and b paths must be in separate XP array clusters. The m and mb paths must be in separate array clusters. The p and m paths must be in separate XP array clusters.</li> <li>• The 2 Gb FCSAs (IOAME) are supported with 1 Gb/2 Gb CHiPs for XP128/1024/10000/12000/20000/24000 and with 4 Gb CHiPs for XP10000/12000/20000/24000.</li> </ul>

**Table 95 NonStop server configuration rules** *(continued)*

Rule number	Description
	<ul style="list-style-type: none"> <li>• The 2 Gb Fibre Channel PICs (VIO) are supported with 1 Gb/2 Gb CHiPs for XP128/1024/10000/12000/20000/24000 and with 4 Gb CHiPs for XP10000/12000/20000/24000.</li> <li>• The 4 Gb Fibre Channel HBAs (in CLIMs) are supported with 1 Gb/2 Gb CHiPs for XP128/1024/10000/12000/20000/24000 and with 4 Gb CHiPs for XP10000/12000/20000/24000.</li> </ul>
15	<p><b>High-availability SAN</b></p> <ul style="list-style-type: none"> <li>• Requires dual-redundant SAN fabrics (level 4, NSPOF high-availability SAN configuration). For information about data availability levels, see “Data availability” (page 38) . Each fabric consists of either a single switch or two cascaded switches, as shown in Figure 64 (page 177) through Figure 71 (page 181).</li> <li>• A single fabric supports a maximum of three switches.</li> <li>• Requires separate fabric zones. Each zone consists of the set of NonStop host (FCSA, Fibre Channel PIC, or CLIM) WWNs and XP storage system port WWNs to be accessed from a single NonStop system.</li> <li>• Configure WWN-based zoning only.</li> <li>• Only NonStop homogeneous connections are allowed to the same zone.</li> <li>• Heterogeneous operating systems can share the same switch or SAN if they are in different zones.</li> <li>• HP recommends host-based mirroring. For example, each LDEV (P) is mirrored to a separate LDEV (M) on separate XP ports (p, b, m, mb paths are used). A nonmirrored volume is allowed. For example, each LDEV (P) is not mirrored to a separate LDEV (M) on separate XP ports (only p and b paths are used).</li> <li>• For high availability, primary (P) LDEVs and mirror (M) LDEVs must be configured on separate array ACP pairs.</li> <li>• For high availability, the p and b paths must be in separate XP array clusters. The m and mb paths must be in separate array clusters.</li> <li>• HP recommends that the p and mb paths be in the same XP array cluster, and the b and m paths be together in the other XP array cluster for a volume.</li> <li>• FCSAs (IOAMEs), Fibre Channel PICs (VIOs), FC HBAs (CLIMs), C-series switches, and B-series switches are supported with 1 Gb/2 Gb CHiPs for XP128/1024/10000/12000/20000/24000 and with 4 Gb CHiPs for XP10000/12000/20000/24000.</li> </ul>

Figure 61 (page 176) shows a minimum direct host attach configuration with an IOAME.

**Figure 61 Minimum direct host attach IOAME configuration for XP storage systems**




Figure 62 (page 176) shows a minimum direct host attach configuration with VIO enclosures.

**Figure 62 Minimum direct host attach VIO configuration for XP storage systems (NS1000, NS14000)**

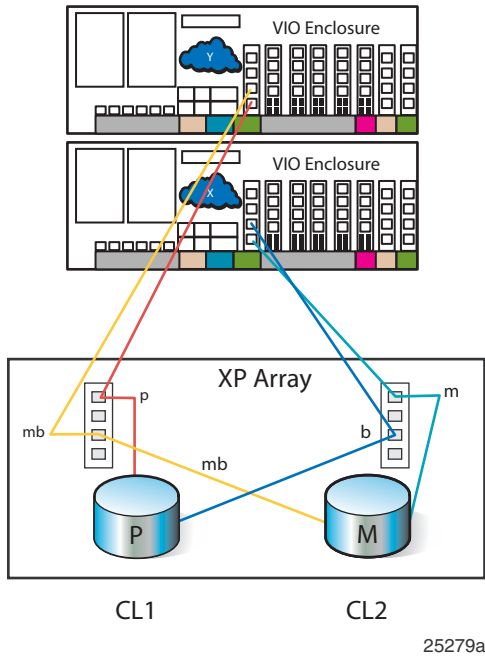





Figure 63 (page 177) shows a minimum direct host attach configuration with CLIMs.


**Figure 63 Minimum direct host attach CLIM configuration for XP storage systems**



26489a

Figure 64 (page 177) shows a minimum SAN configuration with an IOAME.

**Figure 64 Minimum SAN IOAME configuration for XP storage systems**



25123a

Figure 65 (page 178) shows a minimum SAN configuration with VIO enclosures.

**Figure 65 Minimum SAN VIO configuration for XP storage systems (NS1000, NS14000)**




Figure 66 (page 178) shows a minimum SAN configuration with CLIMs.

**Figure 66 Minimum SAN CLIM configuration for XP storage systems**

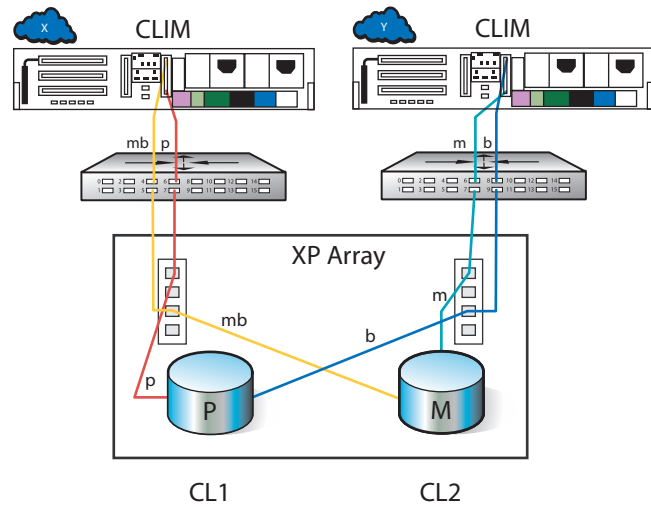


Figure 67 (page 179) shows a configuration with physical IOAME redundancy.

**Figure 67 SAN IOAME configuration with logical and physical redundancy for XP storage systems**

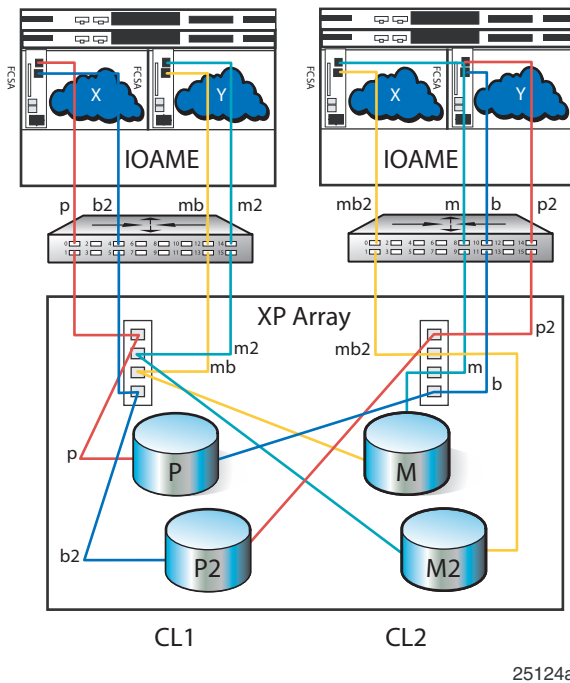


Figure 68 (page 179) shows a SAN configuration with VIO Fibre Channel PIC redundancy.

**Figure 68 SAN VIO configuration with logical and physical redundancy for XP storage systems (NS1000, NS14000)**





Figure 69 (page 180) shows a SAN configuration with CLIM physical redundancy.

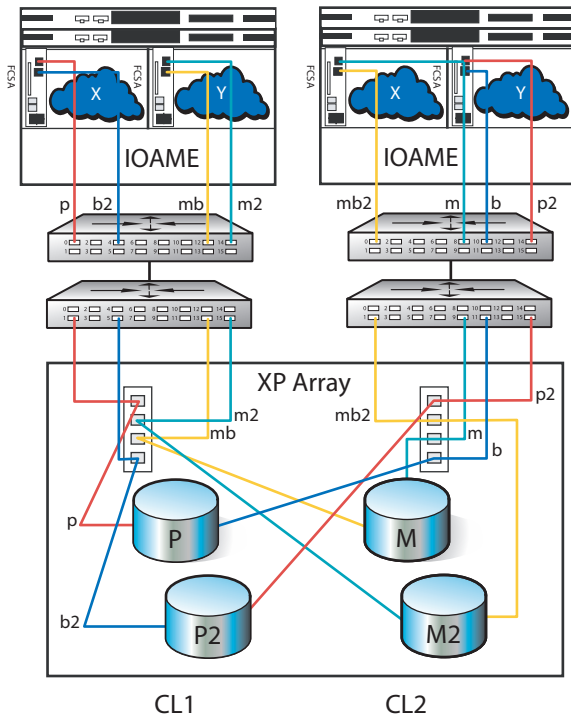
**Figure 69 SAN CLIM configuration with logical and physical redundancy for XP storage systems**



26491a

Figure 70 (page 180) shows a configuration with physical IOAME redundancy.


**Figure 70 SAN IOAME configuration (two cascaded switches) with logical and physical redundancy for XP storage systems**



25125a

Figure 71 (page 181) shows a SAN configuration (two cascaded switches) with VIO Fibre Channel PIC redundancy.


**Figure 71 SAN VIO configuration (two cascaded switches) with logical and physical redundancy for XP storage systems (NS1000, NS14000)**



25278a

Figure 72 (page 182) shows a SAN (two cascaded switches) configuration with CLIM physical redundancy.

**Figure 72 SAN CLIM configuration (two cascaded switches) with logical and physical redundancy for XP storage systems**



26492a

## HP-UX SAN rules

This section describes the SAN rules for HP-UX.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 96 (page 183) describes the SAN configuration rules for HP-UX. Table 97 (page 184) describes support for HP-UX multipathing coexistence.

**Table 96 HP-UX SAN configuration rules**

Storage systems <sup>1</sup>	HP-UX SAN rules
<p>All supported</p> <p><b>Note:</b></p> <p>For 3PAR storage systems, see the 3PAR row in this table.</p>	<ul style="list-style-type: none"> <li>• Supports HP Serviceguard Clusters.</li> <li>• Zoning is required when HP-UX is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports boot from SAN. See “P6000/EVA SAN boot support” (page 221) and “P9000/XP and VA SAN boot support” (page 230).</li> <li>• Supports connection to a common server for mixed storage system types. See “Common SAN coexistence” (page 200).</li> <li>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths. See “Data availability” (page 38).</li> </ul>
<p>P2000 G3 FC MSA2000fc G2 (MSA2300fc)</p>	<ul style="list-style-type: none"> <li>• Host name profile must be set to HP-UX</li> <li>• MSA2012 and MSA2212 FC are not supported with HP-UX</li> <li>• Support for HP-UX 11i v2 (June 2008 or later) using PvLinks</li> <li>• Support for HP-UX 11i v3 using native multipathing</li> </ul>
<p>EVA4000 EVA4100 EVA4400 EVA6000 EVA6100 EVA6400</p> <p>EVA8000 EVA8100 EVA8400 P6300 EVA P6500 EVA</p>	<ul style="list-style-type: none"> <li>• Active/active failover mode is supported for HP-UX 11i v1 and 11i v2 using the Secure Path, PvLinks, or Veritas DMP driver. HP-UX 11i v3 is supported with native multipathing.</li> <li>• For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
<p>EVA3000 EVA5000 VCS 4x</p>	<ul style="list-style-type: none"> <li>• Supports connection of single HBA servers to EVA storage systems.</li> <li>• Active/active failover mode and active/passive failover mode are supported for HP-UX 11.0, 11i v1, and 11i v2 using the Secure Path, PvLinks, or Veritas DMP driver. HP-UX 11i v3 is supported with native multipathing.</li> </ul>
<p>SVSP</p>	<ul style="list-style-type: none"> <li>• Host name profile must be set to HP-UX</li> <li>• Support for HP-UX 11iv2 using Secure Path (legacy only)</li> <li>• Support for HP-UX 11iv3 using native multipathing</li> </ul>
<p>SVSP 3.0</p>	<ul style="list-style-type: none"> <li>• Host name profile must be set to HP-UX</li> <li>• Support for HP-UX 11iv3 using native multipathing</li> <li>• Zoning of host initiator ports, DPM target ports, DPM initiator ports, VSM ports and storage target ports per SVSP requirements</li> </ul>
<p>3PAR F-Class, T-Class</p>	<ul style="list-style-type: none"> <li>• Supported with HP-UX 11iv2 and 11iv3</li> <li>• Supports boot from SAN. See “HP 3PAR SAN boot support” (page 246).</li> <li>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR HP-UX Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 97 HP-UX multipathing coexistence support<sup>1</sup>**

Notes <sup>2</sup>		P2000 G3 MSA2000fc G2 (MSA2300fc)		EVA3000/5000 VCS 4x			P6300/P6500 EVA <sup>4</sup> XCS 1000000 EVA 4000/4100/4400/ 6000/6100/ 8000/8100 XCS 6x <sup>3</sup> 4400 XCS 09x 6400/8400 XCS 095x			XP128/1024/ 10000/12000/ 20000/24000			SVSP 3.0		
		PL	VD	SP	PL	VD	SP	PL	VD	SP	PL	VD	SP	PL	VD
P2000 G3 MSA2000fc G2 (MSA2300fc)	PL	S	—	S	S	—	S	S	—	S	S	—	—	S	—
	VD	—	S	—	—	S	—	—	S	—	—	S	—	—	—
EVA3000/5000 VCS 4x	SP	S	—	S	S	—	S	S	—	S	S	—	—	—	—
	PL	S	—	S	S	—	S	S	—	S	S	—	—	S	—
	VD	—	S	—	—	S	—	—	S	—	—	S	—	—	—
EVA4000/4100/ 4400/6000/6100/ 8000/8100 XCS 6x 4400 XCS 09x 6400/8400 XCS 095x P6300/P6500 EVA <sup>4</sup> XCS 1000000	SP <sup>4</sup>	S	—	S	S	—	S	S	—	S	S	—	—	—	—
	PL	S	—	S	S	—	S	S	—	S	S	—	—	S	—
	VD	—	S	—	—	S	—	—	S	—	—	S	—	—	—
XP128/1024/10000/ 12000/20000/24000	SP	S	—	S	S	—	S	S	—	S	S	—	—	—	—
	PL	S	—	S	S	—	S	S	—	S	S	—	—	S	—
	VD	—	S	—	—	S	—	—	S	—	—	S	—	—	—
SVSP 3.0	SP	—	—	—	—	—	—	—	—	—	—	—	—	—	—
	PL	S	—	—	S	—	—	S	—	—	S	—	—	S	—
	VD	—	—	—	—	—	—	—	—	—	—	—	—	—	—

<sup>1</sup> **Legend:** D = same server and different HBA; S = same server and HBA; — = not supported; sc = single controller; SP = Secure Path; PL = PvLinks; VD = Veritas DMP

<sup>2</sup> General rules:

- P6000/EVA and XP require Secure Path 3.0f SP4 (or later).
- When using Secure Path and Veritas DMP, coexistence is not supported.
- HP-UX 11i v3 is supported with native multipathing. Secure Path is not required or supported. Migration from Secure Path to native multipathing is included in HP-UX 11i v3. For more information, contact an HP storage representative .
- All HP-UX 11i v3 supported storage systems are supported on the same server and HBA using native multipathing software.

<sup>3</sup> EVA4100/6100/8100 requires a minimum of XCS firmware 6.110.

<sup>4</sup> P6300 EVAs and P6500 EVAs do not support Secure Path.

**NOTE:** Secure Path for HP-UX 11.23 is no longer provided through HP. The installations must already have Secure Path for 11.23 already installed.

For more information about storage system coexistence, see “[Heterogeneous SAN storage system coexistence](#)” (page 200).



# HP OpenVMS SAN rules

This section describes the SAN rules for HP OpenVMS.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 98 (page 185) describes the SAN configuration rules for HP OpenVMS. Table 99 (page 186) describes support for HP OpenVMS multipathing coexistence.

**Table 98 HP OpenVMS SAN configuration rules**

Storage systems <sup>1</sup>	OpenVMS SAN rules
All supported	<ul style="list-style-type: none"> <li>• Zoning is required when OpenVMS is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports OpenVMS Clusters.</li> <li>• Supports active/active and active/passive failover mode. A multipathing driver is embedded in the operating system.</li> <li>• Supports boot from SAN. For more information, see “P6000/EVA SAN boot support” (page 221) and “P9000/XP and VA SAN boot support” (page 230).</li> <li>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
EVA3000 EVA6100 EVA4000 EVA6400 EVA4100 EVA8000 EVA4400 EVA8100 EVA5000 EVA8400 EVA6000 P6300 EVA P6500 EVA	<ul style="list-style-type: none"> <li>• For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
P2000 G3 FC P2000 G3 FC/iSCSI MSA2000fc G2 (MSA2300fc)	<ul style="list-style-type: none"> <li>• Host name profile must be set to OpenVMS.</li> <li>• OpenVMS uid must be set on all volumes mapped to a OpenVMS host.</li> <li>• LUN ID 0 is not available for OpenVMS hosts.</li> </ul>
MSA1000	<ul style="list-style-type: none"> <li>• Requires a dedicated MSA1000.</li> <li>• Supports standalone servers or clusters with a maximum of 8 nodes.</li> <li>• Contact an HP storage representative for required operating system patches.</li> </ul>
SVSP	<ul style="list-style-type: none"> <li>• SVSP 3.x is supported by OpenVMS IA54 8.3-1H1, host profile name must be set to OpenVMS.</li> <li>• The number of hosts supported in a OpenVMS cluster is determined by the number of units times the number of hosts times the number of paths &lt; 4096.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 99 HP OpenVMS multipathing coexistence support<sup>1</sup>**

							P6300/P6500 EVA XCS 10000000 EVA 4000/4100/ 6000/6100/ 8000/8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x	
	MSA2000fc G2 (MSA2300fc)	MSA1500 A/A	MSA1000 single controller	EVA 3000/5000 VCS 3x	EVA 3000/ 5000 VCS 4x			XP
Native multipathing driver								
P2000 G3 FC P2000 G3 FC/iSCSI MSA2000fc G2 (MSA2300fc)	S	S	S	S	S	S	S	S
MSA1500 A/A	S	S	S	S	S	S	S	S
MSA1000 single controller	S	S	S	S	S	S	S	S
EVA3000/5000 VCS 3x	S	S	S	S	S	S	S	S
EVA3000/5000 VCS 4x	S	S	S	S	S	S	S	S
EVA4000/ 4100/6000/ 6100/8000/ 8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x P6300/P6500 EVA XCS 10000000	S	S	S	S	S	S	S	S
XP	S	S	S	S	S	S	S	S

<sup>1</sup> **Legend:** D = same server and different HBA; S = same server and HBA; — = not supported

<sup>2</sup> EVA4100/6100/8100 requires XCS firmware 6.110 minimum.

For more information about storage system coexistence, see [“Heterogeneous SAN storage system coexistence”](#) (page 200).

## Host-based volume shadowing

HP OpenVMS servers with host-based volume shadowing are supported in a heterogeneous SAN. Support includes configurations that use host-based volume shadowing over Fibre Channel links with long-distance transceivers or WDM and use of B-series FC or FCIP Fastwrite, or C-series Write

Acceleration. Host-based volume shadowing is supported on the same link with the following applications:

- HP P6000 Continuous Access
- HP P9000 (XP) Continuous Access

**NOTE:** SVSP does not support host-based volume shadowing.

## HP Tru64 UNIX SAN rules

This section describes the SAN rules for HP Tru64 UNIX.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 100 (page 187) describes the SAN configuration rules for HP Tru64 UNIX. Table 101 (page 187) describes support for HP Tru64 UNIX multipathing coexistence.

**Table 100 Tru64 UNIX SAN configuration rules**

Storage systems <sup>1</sup>	Tru64 UNIX SAN rules
All supported	<ul style="list-style-type: none"> <li>• Zoning is required when Tru64 UNIX is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports TruCluster Server.</li> <li>• Supports boot from SAN. For more information, see “P6000/EVA SAN boot support” (page 221) and “P9000/XP and VA SAN boot support” (page 230).</li> <li>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
EVA3000 EVA6000 EVA4000 EVA6100 EVA4100 EVA8000 EVA5000 EVA8100	<ul style="list-style-type: none"> <li>• Zoning is required when a SAN is configured for multiple TruCluster products. Each TruCluster must be in its own zone.</li> <li>• Supports connection of single HBA servers.</li> <li>• Supports active/active or active/passive failover mode.</li> <li>• A multipathing driver is embedded in the operating system.</li> <li>• For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
MSA1000	<ul style="list-style-type: none"> <li>• Requires a dedicated MSA1000.</li> <li>• Supports clusters with a maximum of 4 nodes.</li> <li>• Contact an HP storage representative for required operating system patches.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 101 HP Tru64 UNIX multipathing coexistence support<sup>1</sup>**

		MSA1500 A/A	MSA1000 single controller	EVA3000/ 5000 VCS 3x	EVA3000/ 5000 VCS 4x	EVA 4000/4100/ 6000/6100/ 8000/8100 XCS 6x <sup>2</sup>	XP
		Native multipathing driver					
MSA1500 A/A	Native multipathing driver	S	S	S	S	S	S
MSA1000 single controller		S	S	S	S	S	S
EVA3000/5000 VCS 3x		S	S	S	S	S	S

**Table 101 HP Tru64 UNIX multipathing coexistence support** (continued)

		MSA1500 A/A	MSA1000 single controller	EVA3000/ 5000 VCS 3x	EVA3000/ 5000 VCS 4x	EVA 4000/4100/ 6000/6100/ 8000/8100 XCS 6x <sup>2</sup>	XP
		Native multipathing driver					
EVA3000/5000 VCS 4x		S	S	S	S	S	S
EVA4000/4100/6000 /6100/8000/8100 XCS 6x <sup>2</sup>		S	S	S	S	S	S
XP		S	S	S	S	S	S

<sup>1</sup> Legend: S = same server and HBA

<sup>2</sup> EVA4100/6100/8100 requires XCS firmware 6.cx (or later).

For more information about storage system coexistence, see “Heterogeneous SAN storage system coexistence” (page 200).

## Apple Mac OS X SAN rules

This section describes the SAN rules for Apple Mac OS X.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 102 (page 188) describes the SAN configuration rules for Apple Mac OS X.

**Table 102 Apple Mac OS X SAN configuration rules**

Storage systems	Apple Mac OS X SAN rules
EVA4000	<ul style="list-style-type: none"> <li>Zoning is required when Mac OS X is used in a heterogeneous SAN with other operating systems.</li> <li>Supports active/active and active/passive failover mode. A multipathing driver is embedded in the host adapter driver.</li> <li>Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
EVA4100	
EVA4400	
EVA6000	
EVA6100	
EVA6400	
EVA8000	<p><b>P6000/EVA rules<sup>1</sup></b></p> <ul style="list-style-type: none"> <li>Command View host entry operating system: Apple Mac OS X</li> </ul>
EVA8100	<p><b>Note:</b></p>
EVA8400	

**Table 102 Apple Mac OS X SAN configuration rules (continued)**

Storage systems	Apple Mac OS X SAN rules
<p>P6300 EVA P6500 EVA</p>	<p>For XCS 6.100 (or later) and for XCS earlier than 09003000, select <b>Custom</b> from the Operating System Type list, and then enter <b>00000002024000A8</b> as the Custom type.</p> <p><b>Multipathing rules</b></p> <ul style="list-style-type: none"> <li>• HP recommends zoning single Host Adapter port WWPN with single array port WWPN.</li> <li>• Single Host Adapter port WWPN zoned with multiple array port WWPNs to the same array is supported.</li> <li>• Single Host Adapter port WWPN zoned with multiple arrays is supported.</li> <li>• I/O load balancing can span multiple Host Adapter ports.</li> </ul> <p><b>P6000/EVA rules<sup>1</sup></b></p> <ul style="list-style-type: none"> <li>• Celerity FC-82EN, dual channel, x8 PCI Express</li> <li>• Celerity FC-84EN, quad channel, x8 PCI Express</li> <li>• XCS 6.100 (or later) (EVA4000/4100/6000/6100/8000/8100)</li> <li>• XCS 09000000 (or later) (EVA4400)</li> <li>• Command View 6.0.2 (or later) (EVA4000/4100/6000/6100/8000/8100)</li> <li>• Command View 8.0 (or later) (EVA4400)</li> <li>• Command View host entry operating system: Custom, custom type "00000002024000A8"</li> <li>• HP P6000 Command View 9.4 (or later) (P6300/P6500)</li> </ul> <p><b>Notes:</b></p> <p>For XCS 6.100 (or later), select <b>Custom</b> from the Operating System Type list, and then enter <b>00000002024000A8</b> as the Custom type.</p> <p>For XCS 09003000 (or later), select <b>Apple Mac OS X</b> from the Operating System Type list.</p> <p><b>Host adapter rules</b></p> <ul style="list-style-type: none"> <li>• Supported only with ATTO 8 Gb and 4 Gb PCIe and PCI-x host adapters: <ul style="list-style-type: none"> <li>◦ Celerity FC-41XS, single channel, 64-bit/133 MHz PCI-X</li> <li>◦ Celerity FC-42XS, dual channel, 64-bit/133 MHz PCI-X</li> <li>◦ Celerity FC-41EL, single channel, x4 PCI Express</li> <li>◦ Celerity FC-41ES, single channel, x4 PCI Express</li> <li>◦ Celerity FC-42ES, dual channel, x4 PCI Express</li> <li>◦ Celerity FC-44ES, quad channel, x8 PCI Express</li> </ul> </li> </ul> <p>For more information, see <a href="http://www.attotech.com/hp">http://www.attotech.com/hp</a>.</p>
<p>P2000 G3 FC</p>	<ul style="list-style-type: none"> <li>• Zoning is required when Apple OS X is used in a heterogeneous SAN with other operating systems.</li> <li>• Supported only with the following host adapters: <ul style="list-style-type: none"> <li>◦ Celerity FC-82EN, dual channel, x8 PCI Express</li> <li>◦ Celerity FC-84EN, quad channel, x8 PCI Express</li> </ul> </li> </ul>

<sup>1</sup> Contact an HP storage representative for the latest information on product or firmware version support.

## IBM AIX SAN rules

This section describes the SAN rules for IBM AIX.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 103 (page 190) describes the SAN configuration rules for IBM AIX. Table 104 (page 191) describes support for IBM AIX multipathing coexistence.

**Table 103 IBM AIX SAN configuration rules**

Storage systems <sup>1</sup>	IBM AIX SAN rules
All supported <b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.	<ul style="list-style-type: none"> <li>• Supports HACMP/ES Clusters. For specific version support, contact an HP storage representative.</li> <li>• Zoning is required when IBM AIX is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports active/active and active/passive failover mode. A multipathing driver is required.</li> <li>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths.</li> <li>• For specific Secure Path zoning requirements, contact an HP storage representative.</li> <li>• Supports boot from SAN for XP storage systems. For more information, see “P9000/XP and VA SAN boot support” (page 230).</li> <li>• For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
SVSP	<ul style="list-style-type: none"> <li>• SVSP is supported with AIX 5.3 and 6.1.</li> <li>• Supports HACMP/ES Clusters. For specific version support, contact an HP storage representative.</li> <li>• Zoning is required when IBM AIX is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports active/passive failover mode. A multipathing driver is required.</li> <li>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths (must be Director Class to have dual path in the event of a component failure).</li> <li>• Zoning of host initiator ports, DPM target ports, DPM initiator ports, VSM ports, and storage target ports per SVSP requirements.</li> </ul>
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR AIX Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 104 IBM AIX multipathing coexistence support<sup>1</sup>**

Notes <sup>3</sup>		EVA3000/5000 VCS 3x		EVA 3000/5000 VCS 4x		P6300/P6500 EVA XCS 10000000EVA 4000/4100/6000/ 6100/8000/8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x		XP		SVSP 3.0	
		Secure Path <sup>4</sup>	Antemeta	MPIO	MPIO <sup>5</sup>	Antemeta	HDLM	MPIO	MPIO		
EVA3000/5000 VCS 3x	Secure Path <sup>4</sup>	S <sup>8</sup>	—	D <sup>6</sup>	D <sup>6</sup>	—	—	D <sup>6</sup>	—		
	Antemeta	—	S <sup>7, 8</sup>	—	—	—	—	—	—		
EVA3000/5000 VCS 4x	MPIO	D <sup>6</sup>	—	S <sup>8</sup>	S <sup>8</sup>	—	—	S <sup>8</sup>	S <sup>8</sup>		
EVA4000/4100/ 6000/6100/ 8000/8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x P6300/P6500 EVA XCS 10000000	MPIO <sup>5</sup>	D <sup>6</sup>	—	S <sup>8</sup>	S <sup>8</sup>	—	—	S <sup>8</sup>	S <sup>8</sup>		
	Antemeta	—	—	—	—	S <sup>7, 8</sup>	—	—	—		
XP	HDLM	—	—	—	—	—	S <sup>8</sup>	—	—		
	MPIO	D <sup>6</sup>	—	S <sup>8</sup>	S <sup>8</sup>	—	—	S <sup>8</sup>	S <sup>8</sup>		
SVSP 3.0	MPIO	—	—	S <sup>8</sup>	S <sup>8</sup>	—	—	S <sup>8</sup>	S <sup>8</sup>		

<sup>1</sup> **Legend:** S = same server and HBA; D = same server and different HBA; — = not supported

<sup>2</sup> EVA4100/6100/8100 requires XCS firmware 6.110 minimum.

<sup>3</sup> Connection to a common server with different HBAs requires separate HBA zones for XP, VA, and P6000/EVA.

<sup>4</sup> AIX 5.3 or P5 server support requires Secure Path 2.0d SP3.

<sup>5</sup> EVA4400/6400/8400 requires MPIO 1.0.4.1

<sup>6</sup> Secure Path is supported with current Cambex HBAs. MPIO is supported with IBM native HBAs.

<sup>7</sup> Antemeta is supported with IBM native HBAs.

<sup>8</sup> Requires AIX 5.3 (or later).

For more information about storage system coexistence, see “Heterogeneous SAN storage system coexistence” (page 200).

## Linux SAN rules

This section describes the SAN rules for Linux.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 105 (page 192) describes SAN rules for Red Hat Linux and SUSE Linux.

**Table 105 Linux SAN configuration rules**

Storage systems <sup>1</sup>	Linux SAN rules
All supported <b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.	<ul style="list-style-type: none"> <li>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
P2000 G3 FC MSA2000fc G2 (MSA2300fc) MSA2000fc	<ul style="list-style-type: none"> <li>• Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.</li> <li>• For HBA parameter settings, see “MSA2000fc G2 and MSA2000fc storage system rules” (page 204).</li> <li>• Supports boot from SAN, except for MSA2000fc with the latest versions of RHEL.</li> <li>• Supports Device Mapper.</li> </ul>
MSA1000	<ul style="list-style-type: none"> <li>• Single controller and dual controller with active/active failover</li> </ul>
MSA1500	<ul style="list-style-type: none"> <li>• Single controller and dual controller with active/active failover</li> </ul>
XP48 XP512 XP128 VA7100 XP1024 VA7110 XP10000 VA7400 XP12000 VA7410 XP20000 XP24000	<ul style="list-style-type: none"> <li>• Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.</li> <li>• Provides single-path support only.</li> <li>• Supports clusters (32-bit/64-bit).</li> <li>• Supports active/active or active/passives failover mode.</li> <li>• Supports boot from SAN. For more information, see “P9000/XP and VA SAN boot support” (page 230).</li> </ul>
EVA3000 EVA6100 EVA4000 EVA6400 EVA4100 EVA8000 EVA4400 EVA8100 EVA5000 EVA8400 EVA6000 P6300 EVA P6500 EVA	<ul style="list-style-type: none"> <li>• Zoning is required when Linux is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports connection of single HBA servers.</li> <li>• Supports boot from SAN. For more information, see “P6000/EVA SAN boot support” (page 221).</li> <li>• For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
SVSP	<ul style="list-style-type: none"> <li>• Zoning</li> <li>• Active/Passive support only (no device mapper support for Active/Passive array)</li> <li>• Supports boot from SAN. For more information, see “SVSP SAN boot support” (page 240).</li> </ul>
SVSP 3.0	<ul style="list-style-type: none"> <li>• For information about supported Fibre Channel components, see the SPOCK website at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a>. You must sign up for an HP Passport to enable access.</li> <li>• Supports boot from SAN. For more information, see “P6000/EVA SAN boot support” (page 221).</li> <li>• Direct connect not supported.</li> <li>• Requires a QLogic failover driver.</li> <li>• For information about zoning of host initiator ports, DPM target ports, DPM initiator ports, VSM ports, and storage target ports, see SVSP requirements.</li> </ul>
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR Linux Implementation Guide</i>.</li> </ul>



<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

## Linux multipathing coexistence support

HP does not support coexistence of different multipath products on the same server.

## Microsoft Windows SAN rules

This section describes the SAN rules for Microsoft Windows.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 106 (page 193) describes the SAN configuration rules for Microsoft Windows.

Table 107 (page 194) describes support for Microsoft Windows multipathing coexistence.

**Table 106 Microsoft Windows SAN configuration rules**

Storage systems <sup>1</sup>	Windows SAN rules
All supported <b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.	<ul style="list-style-type: none"> <li>Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
MSA1000 MSA1500	<ul style="list-style-type: none"> <li>Zoning is required when Windows is used in a heterogeneous SAN with other operating systems.</li> <li>Supports a maximum of 20 servers, 8 nodes per cluster.</li> <li>MSA1000 single controller only.</li> <li>MSA1500 single controller and dual controller with active/active failover.</li> </ul>
P2000 G3 FC MSA2000fc G2 (MSA2300fc) MSA2000fc	<ul style="list-style-type: none"> <li>Zoning is required when Windows is used in a heterogeneous SAN with other operating systems.</li> <li>Supports a maximum of 64 hosts.</li> <li>Supports MSCS.</li> <li>Windows 2003 supported with HP MPIO FF DSM 2.6.0.5 (or later).</li> <li>Windows 2008 Microsoft MPIO DSM supported.</li> <li>For HBA parameter settings, see “MSA2000fc G2 and MSA2000fc storage system rules” (page 204).</li> </ul>
VA7100 XP128 VA7110 XP1024 VA7400 XP10000 VA7410 XP12000 XP48 XP20000 XP512 XP24000	<ul style="list-style-type: none"> <li>Zoning is required when Windows is used in a heterogeneous SAN with other operating systems.</li> <li>Supports boot from SAN. For more information, see “P9000/XP and VA SAN boot support” (page 230).</li> </ul>
EVA3000 EVA6100 EVA4000 EVA6400 EVA4100 EVA8000 EVA4400 EVA8100 EVA5000 EVA8400 EVA6000 P6300 EVA	<ul style="list-style-type: none"> <li>Zoning is required when Windows is used in a heterogeneous SAN with other operating systems.</li> <li>In an environment with multiple Windows 2000 or Windows 2003 clusters that share a storage array, configure each cluster and its storage in a separate zone. If the same storage array is used by multiple clusters, the storage ports can be in an overlapping zone.</li> <li>Supports active/active or active/passive failover mode. A multipathing driver is required, which is supplied by HP.</li> <li>Supports MSCS and Microsoft Failover Clustering.</li> </ul>

**Table 106 Microsoft Windows SAN configuration rules** (continued)

Storage systems <sup>1</sup>	Windows SAN rules
P6500 EVA	<ul style="list-style-type: none"> <li>Supports boot from SAN. For more information, see “P6000/EVA SAN boot support” (page 221).</li> <li>For HP P6000 Continuous Access configuration information, see “HP P6000 Continuous Access SAN integration” (page 220).</li> </ul>
SVSP 3.0	<ul style="list-style-type: none"> <li>Zoning is required when Windows is used in a heterogeneous SAN with other operating systems.</li> <li>In an environment with multiple Windows clusters that share a storage array, configure each cluster and its storage in a separate zone. If the same storage array is used by multiple clusters, the storage ports can be in an overlapping zone.</li> <li>Supports only active/passive failover mode. A multipathing driver is required.</li> <li>Supports MSCS and Microsoft Failover Clustering.</li> <li>Supports boot from SAN. For more information, see “SVSP SAN boot support” (page 240).</li> </ul>
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>Supports Hyper-V, Windows 2008, and Windows 2003.</li> <li>Supports Microsoft Failover Clustering and MSCS.</li> <li>Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR Microsoft Windows Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 107 Microsoft Windows multipathing coexistence support<sup>1</sup>**

		P2000 G3 MSA2000fc G2 (MSA2300fc) MSA2000fc			EVA 3000/5000 VCS 4x		P6300/P6500 EVA XCS 10000000  EVA 4000/4100/6000/ 6100/8000/8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x		XP128/1024 XP10000/12000 XP20000/24000		SVSP 3.0	
		MS MPIO DSM <sup>3</sup>	HP MPIO FF <sup>4</sup>	MS MPIO DSM	HP MPIO FF	MS MPIO DSM	HP MPIO FF	MS MPIO DSM	HP MPIO FF	MS MPIO DSM	HP MPIO FF	MS MPIO DSM
P2000 G3 MSA2000fc G2 (MSA2300fc) MSA2000fc	MS MPIO DSM	S	S	S	S	S	S	S	S	S	—	—
EVA3000/5000 VCS 4x	HP MPIO FF	S	S	—	S	S	S	S	S	S	S	—
	MS MPIO DSM	S	—	S	S	S	S	S	S	S	—	—
EVA4000/ 4100/6000/ 6100/8000/ 8100 XCS 6x <sup>2</sup> 4400 XCS 09x	HP MPIO FF	S	S	S	S	—	S	S	S	S	S	—
	HP MPIO DSM	S	S	S	—	S	S	S	S	S	—	—

**Table 107 Microsoft Windows multipathing coexistence support** (continued)

		P2000 G3 MSA2000fc G2 (MSA2300fc) MSA2000fc	EVA 3000/5000 VCS 4x		P6300/P6500 EVA XCS 10000000  EVA 4000/4100/6000/ 6100/8000/8100 XCS 6x <sup>2</sup> 4400 XCS 09x 6400/8400 XCS 095x		XP128/1024 XP10000/12000 XP20000/24000		SVSP 3.0	
			MS MPIO DSM <sup>3</sup>	HP MPIO FF <sup>4</sup>	MS MPIO DSM	HP MPIO FF	MS MPIO DSM	HP MPIO FF	MS MPIO DSM	HP MPIO FF
6400/8400 XCS 095x P6300/P6500 EVA XCS 10000000										
XP128/1024 XP10000/ 12000 XP20000/ 24000	HP MPIO FF	S	S	S	S	S	S	—	S	—
	MS MPIO DSM	S	S	S	S	S	—	S	—	—
SVSP 3.0	MS MPIO DSM	—	S	—	S	—	S	—	S	—
	HP MPIO FF	—	—	—	—	—	—	—	—	—

<sup>1</sup> **Legend:** S = same server and HBA; — = not supported

<sup>2</sup> EVA4100/6100/8100 requires a minimum of XCS firmware 6.110.

<sup>3</sup> Microsoft MPIO Device Specific Module

<sup>4</sup> Microsoft MPIO Full Featured DSM

For more information about storage system coexistence, see “Heterogeneous SAN storage system coexistence” (page 200).

## Sun Solaris SAN rules

This section describes the SAN rules for Sun Solaris.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 108 (page 196) describes the SAN configuration rules for Sun Solaris. Table 109 (page 197) describes support for Sun Solaris multipathing coexistence.

**Table 108 Sun Solaris SAN configuration rules**

Storage systems <sup>1</sup>	Solaris SAN rules
<p>All supported</p> <p><b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.</p>	<ul style="list-style-type: none"> <li>• Zoning is required when Solaris is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports Veritas Cluster Server, Sun Cluster, and PRIMECLUSTER. Contact an HP storage representative for version and storage system support. Not all storage system types are supported on all cluster solutions.</li> <li>• Each cluster must be in its own zone.</li> <li>• Supports connection of single HBA servers.</li> <li>• Supports multipathing high-availability configuration implemented in multiple fabrics or in a single fabric with zoned paths.</li> <li>• Requires a multipathing driver for active/active or active/passive failover with two or more paths.</li> </ul>
<p>P2000 G3 MSA2000fc G2 (MSA2300fc)</p>	<ul style="list-style-type: none"> <li>• Zoning is required when Sun Solaris is used in a heterogeneous SAN with other operating systems.</li> <li>• Single path is not supported. MPxIO is required.</li> <li>• Retry count must be set to 10; entry in <code>/kernel/drv/sd.conf</code> must be set as follows: SD-config-list="HP MSA","retries-timeout: 10"; SD-config-list="HP P2000","retries-timeout: 10";</li> </ul> <p><b>Note:</b> There should be six spaces between "HP" and the array model, "MSA" or "P2000."</p>
<p>EVA4000    XP48 EVA4100    XP512 EVA4400    XP128 EVA6000    XP1024 EVA6100    XP10000 EVA6400    XP12000 EVA8000    XP20000 EVA8100    XP24000 EVA8400    VA7100 P6300 EVA    VA7110 P6500 EVA    VA7400</p>	<ul style="list-style-type: none"> <li>• Supports boot from SAN. For more information, see "P6000/EVA SAN boot support" (page 221) and "P9000/XP and VA SAN boot support" (page 230).</li> <li>• Requires MPxIO, Veritas Volume Manager DMP, or HDLM multipathing driver. For version support, contact an HP storage representative. For multipathing coexistence support, see Table 109 (page 197).</li> <li>• For HP P6000 Continuous Access configuration information, see "HP P6000 Continuous Access SAN integration" (page 220).</li> </ul>

**Table 108 Sun Solaris SAN configuration rules** (continued)

Storage systems <sup>1</sup>	Solaris SAN rules
EVA3000/5000 (VCS 4x) VA7410	
EVA3000/5000 (VCS 3x)	<ul style="list-style-type: none"> <li>Requires Secure Path multipathing driver. For version support, contact an HP storage representative.</li> </ul>
SVSP 3.0	<ul style="list-style-type: none"> <li>SVSP supports Sun Solaris and Solaris clusters as of 3.0.2 release. For more information about Solaris support and supported Fibre Channel components, see the SPOCK website at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a>. You must sign up for an HP Passport to enable access.</li> <li>Requires MPxIO</li> <li>For information about zoning of host initiator ports, DPM target ports, DPM initiator ports, VSM ports, and storage target ports, see the SVSP requirements.</li> </ul>
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR Solaris Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

**Table 109 Sun Solaris multipathing coexistence support<sup>1</sup>**

Notes <sup>3</sup>		P2000 G3 MSA 2000fc G2 (2300fc)		EVA 3000/5000 VCS 3x		EVA 3000/5000 VCS 4x		P6300/P6500 EVA XCS 10000000EVA 4000/4100/6000/6100/8000/8100 XCS 6x <sup>2</sup> 6400/8400 XCS 095x		EVA 4400 XCS 09x		XP			VA		SVSP 3.0
		MP	SP	MP	VD	MP	VD	MP	VD	H	MP	VD	MP	VD	MP		
P2000 G3 MSA2000fc G2 (MSA2300fc)	MP	S	D	S	D	S	D	S	D	D	S	D	S	D	S		
EVA3000/5000 VCS 3x	SP	D	S	D	D	D	D	D	D	—	D	D	D	D	—		
EVA3000/5000 VCS 4x	MP	S	D	S	D	S	D	S	D	—	S	D	S	D	S		
	VD	D	D	D	S	D	S	D	S	—	D	S	D	S	—		
EVA4000/4100/6000/6100/8000/8100 XCS 6x <sup>2</sup>	MP	S	D	S	D	S	D	S	D	—	S	D	S	D	S		
	VD	D	D	D	S	D	S	D	S	—	D	S	D	S	—		
EVA4400 XCS 09x 6400/8400 XCS 095x	MP	S	D	S	D	S	D	S	D	—	S	D	—	—	S		
	VD	D	D	D	S	D	S	D	S	—	D	S	—	—	—		

**Table 109 Sun Solaris multipathing coexistence support** (continued)

Notes <sup>3</sup>	P2000 G3 MSA 2000 G2 (2300x)	EVA 3000/5000 VCS 3x	EVA 3000/5000 VCS 4x		P6300/P6500 EVA XCS 10000000EVA 4000/4100/6000/6100/8000/8100 XCS 6x <sup>2</sup> 6400/8400 XCS 095x		EVA 4400 XCS 09x		XP			VA		SVSP 3.0	
			MP	VD	MP	VD	MP	VD	H	MP	VD	MP	VD	MP	
P6300/P6500 EVA XCS 10000000															
XP	H	D	—	—	—	—	—	—	—	S	—	—	—	—	—
	MP	S	D	S	D	S	D	S	D	—	S	D	S	D	S
	VD	D	D	D	S	D	S	D	S	—	D	S	D	S	—
VA	MP	S	D	S	D	S	D	—	—	—	S	D	S	D	S
	VD	D	D	D	S	D	S	—	—	—	D	S	D	S	—
SVSP 3.0	MP	S	—	S	—	S	—	S	—	—	S	—	S	—	S

<sup>1</sup> **Legend:** S = same server and HBA; D = same server and different HBA; — = not supported, MP = MPxIO, SP = Secure Path, VD = Veritas DMP, H = Hitachi Dynamic Link Manager

<sup>2</sup> EVA4100/6100/8100 requires a minimum of XCS firmware 6.110.

<sup>3</sup> General notes:

- Connection to a common server with different HBAs requires separate HBA zones for XP, VA, and P6000/EVA.
- Secure Path support with the same HBA requires 3.0d minimum. If Secure Path is combined with MPxIO on the same server, Sun-branded HBAs are required for MPxIO.

For more information about storage system coexistence, see “Heterogeneous SAN storage system coexistence” (page 200).

## VMware ESX SAN rules

This section describes the SAN rules for VMware ESX.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 110 (page 199) describes the SAN configuration rules for VMware ESX.

**Table 110 VMware ESX SAN configuration rules**

Storage systems <sup>1</sup>	ESX SAN rules
All supported <b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.	Zoning is required when ESX is used in a heterogeneous SAN with other operating systems.
EVA3000 EVA4000 EVA4100      MSA1000      XP48 EVA4400      MSA1500      XP512 EVA5000      MSA2000fc    XP128 EVA6000      G2              XP1024 EVA6100      (MSA2300fc)    XP10000 EVA6400      MSA2000fc    XP12000 EVA8000      P2000 G3      XP20000 EVA8100      FC EVA8400      SVSP            XP24000 P6300 EVA P6500 EVA	Supports Microsoft Windows Server 2003 base, with a maximum of two nodes per cluster. It does not support Windows Server 2003 SP1. <b>Note:</b> Not all listed storage systems are supported on all versions of VMware ESX.
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR VMware ESX 3.x – 4.x Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

## VMware ESX multipathing coexistence support

All supported storage systems are supported on the same server using native multipathing software.

**NOTE:** For 3PAR storage systems, see “Common server, different HBAs” (page 201).

## Citrix Xen SAN rules

This section describes the SAN rules for Xen.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 111 (page 199) describes SAN rules for Xen

**Table 111 Xen SAN configuration rules**

Storage systems <sup>1</sup>	Xen SAN rules
All supported <b>Note:</b> For 3PAR storage systems, see the 3PAR row in this table.	<ul style="list-style-type: none"> <li>• Supports multipathing high-availability configuration in multiple fabrics or in a single fabric with zoned paths.</li> </ul>
P2000 G3 FC MSA2000fc G2 (MSA2300fc)	<ul style="list-style-type: none"> <li>• Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.</li> <li>• For HBA parameter settings, see “MSA2000fc G2 and MSA2000fc storage system rules” (page 204).</li> </ul>

**Table 111 Xen SAN configuration rules** *(continued)*

Storage systems <sup>1</sup>	Xen SAN rules
MSA2000fc	<ul style="list-style-type: none"> <li>• Boot from SAN is not supported.</li> <li>• Supports Inbox Device Mapper.</li> </ul>
XP10000    XP20000 XP12000    XP24000	<ul style="list-style-type: none"> <li>• Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports active/active or active/passives failover mode.</li> <li>• Boot from SAN is not supported.</li> </ul>
EVA4000    EVA8000 EVA4100    EVA8100 EVA4400    EVA8400 EVA6000    P6300 EVA EVA6100    P6500 EVA EVA6400	<ul style="list-style-type: none"> <li>• Zoning is required when Xen is used in a heterogeneous SAN with other operating systems.</li> <li>• Supports connection of single HBA servers.</li> <li>• Boot from SAN is not supported.</li> <li>• For HP P6000 Continuous Access configuration information, see <a href="#">“HP P6000 Continuous Access SAN integration”</a> (page 220).</li> </ul>
3PAR F-Class, T-Class	<ul style="list-style-type: none"> <li>• Zoning by HBA is required when used in a heterogeneous SAN, including other operating systems and other storage system families or types.</li> <li>• All hosts must have the appropriate Host Operating System type parameter set (Host Persona) and the required host settings described in the <i>3PAR Citrix XenServer Implementation Guide</i>.</li> </ul>

<sup>1</sup> Unlisted but supported storage systems have no additional SAN configuration restrictions. For the latest support information, contact an HP storage representative.

## Heterogeneous SAN storage system coexistence

HP supports SANs that contain different storage system types. This section defines rules for mixing storage system types in a SAN.

Any combination of storage systems is supported in a SAN, provided that the configuration conforms to all applicable fabric rules. For more information, see:

- [“B-series switches and fabric rules”](#) (page 83)
- [“C-series switches and fabric rules”](#) (page 115)
- [“H-series switches and fabric rules”](#) (page 131)
- [“M-series switches and fabric rules”](#) (page 140)
- [“Heterogeneous server rules”](#) (page 159)

## Common SAN coexistence

To configure different HP storage system types or third-party storage systems for coexistence in a common SAN, without common access from the same server, define a separate zone for each storage system family. Storage system families include:

- P2000, MSA
- P6000, EVA
- SVSP
- 3PAR

---

**NOTE:** For Fibre Channel fabric support for 3PAR storage systems, see [“HP 3PAR storage system rules”](#) (page 241).

---

- P9500, XP, VA
- Third-party storage systems



To configure common server access to different storage system types, see [“Common server access” \(page 201\)](#).

## Common server access

Common server access enables simultaneous connectivity to different storage system types from the same server. In some cases, connectivity to different storage system types, including disk and tape, is allowed from the same HBA.

For supported common server access configurations and rules, see the storage solutions documentation.

Common server access to HP storage and third-party storage is supported if you comply with the following rules:

- You must use HP-supported HBAs, parameter settings, drivers, and firmware to support HP storage and third-party storage on the same HP or multi-vendor server.
- Use different HP-supported HBA vendors if the required parameter settings, drivers, and firmware are different for HP storage and the third-party storage.

## Common server, different HBAs

For information about supported storage system types using common server access with different HBAs (Emulex and QLogic only), see the multipathing coexistence table (in the section for your operating system section).

---

**NOTE:** HP Brocade-branded HBAs are not supported for coexistence with other vendor HBAs.

---

For all multipath configurations, all paths for a given LUN must only be accessed by HBAs from the same vendor (for example, Emulex or QLogic, but not both).

- IBM AIX ([“IBM AIX multipathing coexistence support” \(page 191\)](#))
- Sun Solaris ([“Sun Solaris multipathing coexistence support” \(page 197\)](#))

For all other operating systems, connection to a common server with different HBA vendor products requires separate HBA zones for each storage system:

- All Fibre Channel HBA zones must contain HBAs from the same vendor.
- A zone can contain different HBA models if they are all from the same HBA vendor.
- A Fibre Channel HBA can be a member in more than one zone.
- All HBA members in the same zone can reside in different servers, but must be the same operating system type.

Failure to properly zone different HBA models may result in lost connectivity and instability of the operating system.

Before implementation, contact an HP storage representative for a configuration review.

---

**NOTE:** For zoning requirements and recommendations on server operating system multipathing settings, SCSI timeout settings, and HBA driver settings on 3PAR storage systems, see the 3PAR implementation guides for each of the server operating systems.

---

## Common server, common HBA

Simultaneous access to different storage system types from the same HBA is supported for some storage systems. Each storage system must support identical:

- HBA model numbers
- Driver versions
- Multipathing software versions or compatible multipathing products

See the multipathing coexistence tables (in each operating system's section) for supported storage system types using common server access with a common HBA:

- [“HP-UX multipathing coexistence support” \(page 184\)](#)
- [“HP OpenVMS multipathing coexistence support” \(page 186\)](#)
- [“HP Tru64 UNIX multipathing coexistence support” \(page 187\)](#)
- [“IBM AIX multipathing coexistence support” \(page 191\)](#)
- [“Microsoft Windows multipathing coexistence support” \(page 194\)](#)
- [“Sun Solaris multipathing coexistence support” \(page 197\)](#)

---

**NOTE:** 3PAR storage systems are not supported for common HBA access with other storage system families. For information about configuring different storage system families on the same server with different HBAs, see [“Common server, different HBAs” \(page 201\)](#). For information about data migration between 3PAR storage systems and other storage system families, see [“HP 3PAR data migration” \(page 243\)](#).

---

## Server zoning rules

A SAN configuration must conform to the zoning requirements of the operating system and other SAN components. For specific information, see:

- [H-series—“H-series switch zoning limits” \(page 138\) and “Zoning enforcement for H-series switches” \(page 139\)](#)
- [B-series—“Zoning enforcement for B-series Fibre Channel switches and MP Router LANS” \(page 113\)](#)
- [C-series—“Zoning limits for C-series Fibre Channel switches ” \(page 129\) and “Zoning enforcement for C-series Fibre Channel switches” \(page 129\)](#)
- [M-series—Table 69 \(page 144\) and “Zoning enforcement for M-series and McDATA switches” \(page 145\)](#)
- [“Zoning” \(page 208\) for MSA1500](#)
- [“Zoning” \(page 211\) for MSA1000](#)
- [“Configuration parameters” \(page 216\) for P6000/EVA](#)
- [“Zoning” \(page 221\) for P6000/EVA](#)
- [“Zoning” \(page 229\) for XP and VA](#)
- [“Zoning enforcement” \(page 376\)](#)
- [“Zoning guidelines” \(page 376\)](#)
- [“EBS zoning” \(page 377\)](#)
- [“Zone naming” \(page 377\)](#)

# 11 P2000 and MSA storage system rules

This chapter describes specific rules for the following entry-level storage systems:

- “HP StorageWorks P2000 G3 FC/iSCSI, Modular Smart Array 2000fc G2, and Modular Smart Array 2000fc” (page 203)
- “HP StorageWorks Modular Smart Array 1500” (page 207)
- “HP StorageWorks Modular Smart Array 1000” (page 209)

For the iSCSI rules for the MSA2000i and MSA1510i, respectively, see “HP StorageWorks MSA/2000 family of iSCSI SAN arrays” (page 318) and “HP StorageWorks Modular Smart Array 1510i” (page 317).

## HP StorageWorks P2000 G3 FC/iSCSI, Modular Smart Array 2000fc G2, and Modular Smart Array 2000fc

Table 112 (page 203) describes the P2000 G3 FC/iSCSI and MSA2000fc G2 chassis configurations.

**Table 112 P2000 G3 FC/iSCSI, MSA2000fc G2, and MSA2000fc chassis configurations**

Model	Array chassis	Expansion options <sup>1</sup>	Expansion maximum
P2000 FC P2000 G3 FC/iSCSI	P2000 3.5-in drive bay (LFF) P2000 2.5-in drive bay (SFF) 2012 3.5-in drive bay (LFF, upgrade only) 2024 2.5-in drive bay (SFF, upgrade only)	P2000 LFF drive enclosure D2700 SFF disk enclosure MSA2000 LFF drive enclosure MSA70 SFF drive enclosure	8 total enclosures or 149 drives 96 LFF drives 149 SFF drives
MSA2000fc G2 (MSA2312fc/MSA2324fc)	2012 3.5-in drive bay (LFF) 2024 2.5-in drive bay (SFF)	MSA2000 LFF drive enclosure MSA70 SFF drive enclosure	5 total enclosures or 99 drives 60 LFF drives 99 drives
MSA2000fc (MSA2012fc/MSA2212fc)	Pre-configured single or dual controller 3.5-in drive bay (LFF)	MSA2000 LFF drive enclosure	4 total enclosures 48 LFF drives

<sup>1</sup> Confirm cabling requirements and limitation in the cable configuration guide at <http://www.hp.com/go/msa>.

Table 113 (page 203) describes the MSA2000fc controller configurations.

**Table 113 MSA2000fc controller configurations**

Storage system	Description
MSA2012fc single controller	Used in a direct connect or SAN connect configuration with a standard single controller
MSA2012fc dual controller	Used in a direct connect or SAN connect configuration with standard dual controllers
MSA2212fc enhanced dual controller	Used in a direct connect or SAN connect configuration with enhanced dual controllers

## Heterogeneous SAN support

The P2000 G3 and MSA2000 family supports heterogeneous access with any of the operating systems listed in Table 114 (page 204). Standalone servers and clustered servers are supported on the same storage system.

**Table 114 P2000 G3 and MSA2000 family heterogeneous SAN support**

Storage system	Firmware	Switches	Operating systems
P2000 G3 FC P2000 G3 FC/iSCSI	TS230P006 (or later)	H-series <sup>1</sup> B-series C-series M-series	Apple OS X HP-UX OpenVMS Microsoft Windows Red Hat Enterprise Linux Sun Solaris SUSE Linux Enterprise Server VMware ESX
MSA2000fc G2 (MSA2312fc/MSA2324fc)	M112R14-03 (or later)		HP-UX Microsoft Windows OpenVMS Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX Sun Solaris
MSA2000fc (MSA2012fc/MSA2212fc)	J201R09-01 (or later)	H-series B-series C-series	Linux Microsoft Windows VMware ESX

<sup>1</sup> The H-series switches have limited operating system support. See Table 55 (page 135).

For the latest information on firmware versions and MSA storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## Configuration rules

Table 115 (page 204) lists the MSA2000fc and MSA2000fc G2 storage system rules.

**Table 115 MSA2000fc G2 and MSA2000fc storage system rules**

Rule number	Description
1	The MSA2000fc G2, MSA2000fc, and P2000 G3 FC are supported in all SAN fabric topology configurations described in this guide.
2	Unless otherwise specified, the MSA2000fc G2, MSA2000fc, and P2000 G3 FC can be configured in a SAN using the switch models listed in this guide. See: <ul style="list-style-type: none"> <li>• Table 53 (page 132) (H-series switches)</li> <li>• Table 21 (page 85) through Table 23 (page 87) (B-series)</li> <li>• Table 38 (page 118) (C-series)</li> </ul>
3	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.
4	MSA2000fc—When connecting to a SAN, disable Port Interconnect in the HP Storage Management Utility. (See the <i>HP StorageWorks 2000fc Modular Smart Array User Guide</i> .)
5	MSA2000fc—Do not use the MSA2000fc port WWID in your zoning configuration. Port 1 on the controller and the node share a WWID.
6	Linux with Emulex or QLogic—See the <i>Installation and Reference Guide Device Mapper Multipath Enablement Kit for HP StorageWorks Disk Arrays</i> . For all other supported operating systems, set the following parameters:

**Table 115 MSA2000fc G2 and MSA2000fc storage system rules** (continued)

Rule number	Description
	<ul style="list-style-type: none"><li>• <b>Brocade:</b> MPIO Mode Enable = off, Path Timeout = 60</li><li>• <b>Emulex:</b> LinkTimeOut = 60, NodeTimeOut = 60</li><li>• <b>QLogic:</b> Port Down Retry Count = 60, Link Down Timeout = 60</li></ul>
7	MSA2000fc—The primary controller presents the volumes. If the primary controller fails, the secondary controller presents the volumes.
8	Use volume mapping to enable or disable volume access to specific hosts.
9	When connecting to a SAN, HP recommends setting the topology to "Point to Point" on the MSA2000fc G2 and P2000 G3 FC. HP recommends setting the topology to "Loop" in all cases for the MSA2000fc.

## Zoning

These zoning rules apply to a heterogeneous SAN with P2000 G3 FC, FC/iSCSI, MSA2000fc G2, or MSA2000fc storage systems:

- Zoning is required for all operating systems that access the array. See [“Heterogeneous SAN support”](#) (page 203).
- The array can be in multiple operating system zones.

For configurations that support common server access to multiple storage system types, see [“Common server access”](#) (page 201).

## Maximums

[Table 116](#) (page 206) lists maximum configurations for the P2000 G3 FC, FC/iSCSI, MSA2000fc G2, and MSA2000fc.

**Table 116 P2000 G3 FC, FC/iSCSI, MSA2000fc G2, and MSA2000fc maximum configurations**

Storage systems	Operating systems	Drives	Hosts	LUNs	LUN size	Snapshots and clones <sup>1</sup>
P2000 G3 FC P2000 G3 FC/iSCSI	HP-UX 11.23, 11.31 OpenVMS 8.3, 8.3-1H1 Citrix Xen Server HP Enterprise Edition Microsoft Windows Server 2003, 2008, Hyper-V Red Hat Enterprise Linux 4, 5 SUSE Linux Enterprise Server 10, 11 VMware ESX 3.5, 4.0 Sun Solaris 10	96 3.5-in LFF 149 2.5-in SFF	64	512	Up to 64 TB depending on vdisk configuration	64 Standard (maximum 512 Snapshots, clones or remote snaps)
MSA2000fc G2 (MSA2312fc /MSA2324fc)	HP-UX 11.23, 11.31 OpenVMS 8.3, 8.3-1H1 Citrix Xen Server HP Enterprise Edition Microsoft Windows Server 2003, 2008, Hyper-V Red Hat Enterprise Linux 4, 5 SUSE Linux Enterprise Server 10, 11 VMware ESX 3.0, 3.5, 4.0 Sun Solaris 10	60 3.5-in LFF 99 2.5-in SFF	64	512	Up to 64 TB depending on vdisk configuration	Up to 255 snapshots or clones
MSA2000fc (MSA2012fc /MSA2212fc)	Citrix Xen Server HP Enterprise Edition Microsoft Windows Server 2003, 2008, Hyper-V Red Hat Enterprise Linux 4, 5 SUSE Linux Enterprise Server 9, 10 VMware ESX 3.0, 3.5, 4.0 Sun Solaris 10	48	64	256 <sup>2</sup>	16 TB	Up to 64 snapshots Up to 128 volume copies

<sup>1</sup> Snapshots and clones require an additional license.

<sup>2</sup> A single controller supports 128 LUNs. Two controllers are required for the maximum of 256 LUNs.

Contact an HP storage representative for a configuration review to determine support for configurations that exceed the stated limits.

## Management software support

The P2000 G3 FC, FC/iSCSI, MSA2000fc G2, and MSA2000fc support target-based management interfaces, including Telnet (CLI), FTP, and a web-based interface. The web-based interface is supported with Microsoft Internet Explorer 6 (or later) and Mozilla Firefox 1.0.7 (or later).

## Tape storage

The following rules apply to tape storage support in a SAN with P2000 G3 FC, FC/iSCSI, MSA2000fc G2, or MSA2000fc storage systems:

- Overlapping zones are supported with disk and tape.
- Common HBAs for shared disk and tape connections are supported. Shared tape connections can exist in a disk multipath environment, but there is no multipath support for tape connections because they are zoned to only one host HBA. For optimal performance, HP recommends using separate HBAs for disk and tape connections.
- For a VLS, multiple heterogeneous initiators and multiple ports of the VLS device are supported in the same fabric zone.

# HP StorageWorks Modular Smart Array 1500

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions

For additional support information, see *HP StorageWorks Modular Smart Array 1500 cs Compatibility Guide*, available at <http://www.hp.com/go/msa1500cs>. Under Related Information, click **HP Support & Drivers**, select your product, and then click **Manuals**.

## Homogeneous SAN support

A single MSA 1500 supports homogeneous access with the operating systems listed in [Table 117 \(page 207\)](#). Each operating system requires a dedicated MSA 1500. Standalone servers and clustered servers are supported on the same storage system.

**Table 117 MSA1500 homogeneous SAN support**

Storage system	Firmware <sup>1</sup>	Switches	Operating systems <sup>2</sup>
MSA1500 active/active	7x, 6x	B-series	HP OpenVMS (7x firmware only)
		C-series	HP Tru64 (7x firmware only)
		M-series	SCO
MSA1500 active/passive or single controller	5x, 4x	B-series	HP-UX (5x firmware only)
		C-series	VMware ESX (4x firmware only)
		M-series	

<sup>1</sup> For the latest firmware support information, contact an HP storage representative.

<sup>2</sup> For the supported operating system versions, see “[Heterogeneous server rules](#)” (page 159).

For the latest information on firmware versions and MSA storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## Heterogeneous SAN support

A single MSA 1500 supports heterogeneous access with any of the operating systems listed in [Table 118 \(page 207\)](#). Standalone servers and clustered servers are supported on the same storage system.

**Table 118 MSA1500 heterogeneous SAN support**

Storage system	Firmware <sup>1</sup>	Switches	Operating systems <sup>2</sup>
MSA1500 active/active	7x, 6x	B-series	HP-UX
		C-series	Linux
		H-series	Microsoft Windows
		M-series	VMware ESX (7x firmware only)
MSA1500 active/passive or single controller	5x, 4x	B-series	Microsoft Windows (5x firmware only)
		C-series	Linux
			VMware ESX (5x firmware only)

<sup>1</sup> For the latest firmware support information, contact an HP storage representative.

<sup>2</sup> For the supported operating system versions, see “[Heterogeneous server rules](#)” (page 159).

For the latest information on firmware versions and MSA storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

---

**NOTE:** There are heterogeneous restrictions based on failover mode and clustering.

---

## Configuration rules

Table 119 (page 208) lists the MSA1500 storage system rules.

**Table 119 MSA1500 storage system rules**

Rule number	Description
1	The MSA1500 is supported in all SAN fabric topology configurations described in this guide.
2	Unless otherwise specified, the MSA1500 can be configured in a SAN using the switch models listed in this guide. See: <ul style="list-style-type: none"> <li>• Table 53 (page 132) (H-series)</li> <li>• Table 21 (page 85) through Table 23 (page 87) (B-series)</li> <li>• Table 38 (page 118) (C-series)</li> <li>• Table 61 (page 140) and Table 62 (page 141) (M-series)</li> </ul>
3	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.

## Zoning

These zoning rules apply to a heterogeneous SAN with MSA1500 storage systems:

- Zoning is required for all operating systems that access an MSA1500. See “Heterogeneous server rules” (page 159).
- An MSA1500 storage system can be in multiple operating system zones.

For configurations that support common server access to multiple storage system types, see “Common server access” (page 201).

## Maximums

Table 120 (page 208) lists maximum configurations for the MSA1500.

**Table 120 MSA1500 maximum configurations**

Storage system	Operating systems <sup>1</sup>	HBAs per server	Active controller ports (targets) per HBA	LUNs per HBA target	HBAs (initiators) per active controller port	Maximum LUNs per MSA
MSA1500	HP-UX Red Hat Linux SUSE Linux Microsoft Windows 2000 Microsoft Windows 2003 VMware ESX	2	8	32	32	32

<sup>1</sup> For the supported operating system versions, see “Heterogeneous server rules” (page 159).

Contact an HP storage representative for a configuration review to determine support for configurations that exceed the stated limits.



## Tape storage

The following rules apply to tape storage support in a SAN with MSA1500 storage systems:

- Overlapping zones are supported with disk and tape.
- Separate or common HBAs for disk and tape connections are supported.
- For a VLS, multiple heterogeneous initiators and multiple ports of the VLS device are supported in the same fabric zone.
- HP recommends using a separate tape-to-HBA connection for servers with backups that require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your HP representative.

## HP StorageWorks Modular Smart Array 1000

The MSA1000 is available in four product configurations:

- **MSA1000 with optional embedded hub**

MSA1000 storage systems with the MSA Hub 2/3. An MSA1000 using this hub cannot cascade to any other interconnect; it is designed only for a two-node configuration. This configuration is supported with 4 Gb HBAs.

- **MSA1000 standalone storage system**

A 2 Gb Fibre Channel storage system for entry-level and mid-range SANs.

- **MSA1000 with optional embedded switch**

MSA1000 storage systems with the MSA SAN Switch 2/8 are supported with B-series switches only.

- **MSA1000 with Small Business SAN G2 kit**

The HP StorageWorks Modular Smart Array 1000 Small Business SAN G2 kit is a 2 Gb Fibre Channel storage system designed to reduce the complexity and risk of an entry-level SAN. It provides a low-cost, scalable, high-performance storage consolidation system with investment protection.

The MSA1000 Small Business SAN is a standalone SAN configuration that is not supported for connectivity to B-series, C-series, H-series, or M-series SAN fabrics.

For configuration support information, see <http://h18006.www1.hp.com/products/storageworks/msa1000smb/index.html>.

This section describes support for MSA1000 storage systems that are not in a Small Business SAN.

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware version
- HBA firmware and driver versions

For additional support information, see the *MSA1000 Compatibility Guide* under Technical documentation at <http://www.hp.com/go/msa1000>.

## Homogeneous SAN support

The MSA1000 supports homogeneous access with the operating systems listed in [Table 121 \(page 210\)](#). Each operating system requires a dedicated MSA1000. Standalone servers and clustered servers are supported on the same storage system.

**Table 121 MSA1000 homogeneous SAN support**

Storage system	Firmware <sup>1</sup>	Switches	Operating systems <sup>2</sup>
MSA1000 active/active	7x	B-series C-series M-series	HP OpenVMS
			HP Tru64 UNIX
MSA1000 active/passive or single controller	5x, 4x	B-series C-series M-series	HP-UX
			VMware ESX (4x firmware only)

<sup>1</sup> For the latest firmware support information, contact an HP storage representative.

<sup>2</sup> For the supported operating system versions, see “[Heterogeneous server rules](#)” (page 159).

## Heterogeneous SAN support

A single MSA1000 supports heterogeneous access with the operating systems listed in [Table 122](#) (page 210). Standalone servers and clustered servers are supported on the same storage system.

**Table 122 MSA1000 heterogeneous SAN support**

Storage system	Firmware <sup>1</sup>	Switches	Operating systems <sup>2</sup>
MSA1000 active/active	7x	B-series	HP-UX
		C-series	Microsoft Windows
		H-series	Linux
		M-series	VMware ESX
MSA1000 active/passive or single controller	5x, 4x	B-series	Microsoft Windows
		C-series	Linux (4x firmware only)
		M-series	VMware ESX (5x firmware only)

<sup>1</sup> For the latest firmware support information, contact an HP storage representative.

<sup>2</sup> For the supported operating system versions, see “[Heterogeneous server rules](#)” (page 159).

## Configuration rules

[Table 123](#) (page 211) lists the MSA1000 storage system rules.

**Table 123 MSA1000 storage system rules**

Rule number	Description
1	The MSA1000 is supported in all SAN fabric topology configurations described in this guide.
2	Unless otherwise specified, the MSA1000 can be configured in a SAN using the switch models listed in this guide. See: <ul style="list-style-type: none"> <li>• <a href="#">Table 53 (page 132)</a> (H-series)</li> <li>• <a href="#">Table 21 (page 85)</a> through <a href="#">Table 23 (page 87)</a> (B-series)</li> <li>• <a href="#">Table 38 (page 118)</a> (C-series)</li> <li>• <a href="#">Table 61 (page 140)</a> and <a href="#">Table 62 (page 141)</a> (M-series)</li> </ul>
3	Multipathing is supported for the following operating systems: <ul style="list-style-type: none"> <li>• HP-UX</li> <li>• OpenVMS</li> <li>• Tru64 UNIX</li> <li>• Linux</li> <li>• Microsoft Windows</li> <li>• VMware ESX</li> </ul>
4	Attachment of single-path servers to a dual-controller MSA1000 configured with Secure Path and dual HBA servers is supported with Microsoft Windows. However, if the active controller fails, single-path servers lose access to data on the MSA1000.
5	If Secure Path for Linux is used on any node or cluster attached to an MSA1000, all nodes must also have Secure Path installed, regardless of operating system. For the Secure Path parameters, see <a href="http://www.hp.com/go/securepath">http://www.hp.com/go/securepath</a> .
6	To enable or disable LUN access to specific connections, use ACU.

## Zoning

These zoning rules apply to a heterogeneous SAN with MSA1000 storage systems:

- Zoning is required for all operating systems that access an MSA1000. See “[Heterogeneous server rules](#)” (page 159).
- An MSA1000 storage system can be in multiple operating system zones.

For configurations that support common server access to multiple storage system types, see “[Common server access](#)” (page 201).

## Maximums

[Table 124 \(page 212\)](#) lists maximum configurations for the MSA1000.

**Table 124 MSA1000 maximum configurations**

Storage system	Operating systems <sup>1</sup>	HBAs per server	Active controller ports (targets) per HBA	LUNs per HBA target	HBAs (initiators) per active controller port	Maximum LUNs per MSA
MSA1000	HP-UX Red Hat Linux SUSE Linux Microsoft Windows VMware ESX	2	8	32	32	32
	Tru64 UNIX OpenVMS	Server dependent	8	32	32	32

<sup>1</sup> For the supported operating system versions, see “Heterogeneous server rules” (page 159).

Contact an HP storage representative for a configuration review to determine support for configurations that exceed the stated limits.

## Tape storage

The following rules apply to tape storage support in a SAN with MSA1000 storage systems:

- Overlapping zones are supported with disk and tape.
- Separate or common HBAs for disk and tape connections are supported.
- For a VLS, multiple heterogeneous initiators and multiple ports of the VLS device are supported in the same fabric zone.
- HP recommends using a separate tape-to-HBA connection for servers with backups that require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your HP representative.

---

## 12 P6000/EVA storage system rules

This chapter describes specific rules for the following storage systems:

- EVA3000/5000
- EVA4000/4100
- EVA4400
- EVA6000/6100
- EVA8000/8100
- EVA6400/8400
- P6300/P6500 EVA

---

❗ **IMPORTANT:** HP P6000 storage was formerly called the HP Enterprise Virtual Array product family. General references to HP P6000 can also refer to earlier versions of HP EVA products.

---

This chapter describes the following topics:

- “P6000/EVA storage” (page 213)
- “P6000/EVA data migration” (page 217)
- “HP P6000 Continuous Access SAN integration” (page 220)
- “Zoning” (page 221)
- “P6000/EVA SAN boot support” (page 221)
- “Storage management server integration” (page 221)
- “Cabling” (page 222)

For information about the EVA iSCSI Connectivity Option, see “EVA and EVA4400 iSCSI Connectivity Option ” (page 333).

### P6000/EVA storage

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions

## Heterogeneous SAN support

P6000/EVA HSV-based controller storage systems support shared access with any combination of operating systems listed in [Table 125 \(page 214\)](#).

**Table 125 P6000/EVA heterogeneous SAN support**

Storage systems <sup>1</sup>	Firmware <sup>2</sup>	Switches	Operating systems <sup>3</sup>
EVA3000/5000	VCS 4x, 3x	B-series C-series H-series <sup>4</sup> M-series	Apple Mac
EVA4000/4100	XCS 6x		Citrix Xen Server
EVA6000/6100			HP-UX
EVA8000/8100			IBM AIX
EVA4400	XCS 09x XCS 10000000	M-series	Microsoft Windows
EVA6400/8400	XCS 095x XCS 10000000		OpenVMS
			Red Hat Linux
P6300/P6500 EVA	XCS 10000000	B-series C-series H-series <sup>4</sup>	Sun Solaris
			SUSE Linux
			Tru64 UNIX <sup>5</sup>
			VMware ESX
			<b>Note:</b> For the latest information on storage system support, see the SPOCK website at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.

<sup>1</sup> EVA4100/6100/8100 requires XCS firmware 6.2x minimum.

<sup>2</sup> Contact an HP storage representative for switch model and operating system support when using XCS 10x, 09x or 6x.

<sup>3</sup> For the supported operating system versions, see “[Heterogeneous server rules](#)” (page 159).

<sup>4</sup> The H-series switches are not supported with EVA3000/5000 VCS 3x and has limited operating system support. See “[H-series switches servers, operating systems, and storage system support](#)” (page 135).

<sup>5</sup> The P6300/P6500 EVA, EVA6400/8400 and EVA4400 are not supported with Tru64 UNIX.

## Configuration rules

[Table 126 \(page 214\)](#) lists configuration rules for heterogeneous SAN storage with P6000/EVA storage systems.

**Table 126 P6000/EVA storage system rules**

Rule number	Description
1	P6000/EVA storage systems are supported in all SAN fabric topology configurations described in this guide.
2	Unless otherwise specified, P6000/EVA storage systems can be configured in a SAN using the switch models listed in this guide. See the following tables: <ul style="list-style-type: none"> <li>• <a href="#">Table 53 (page 132)</a> (H-series switches)</li> <li>• <a href="#">Table 21 (page 85)</a> through <a href="#">Table 23 (page 87)</a> (B-series)</li> <li>• <a href="#">Table 38 (page 118)</a> (C-series)</li> <li>• <a href="#">Table 61 (page 140)</a> and <a href="#">Table 62 (page 141)</a> (M-series)</li> </ul>
3	For SANs with more than 1,024 HBAs, an HSV controller must be zoned for a maximum of 1,024 HBAs. Add a zone to a SAN to satisfy the 1,024 HBA limit.
4	When using HP P6000 Continuous Access, the zoning limit is 256 HBAs.
5	For shared access and heterogeneous platform zoning requirements, see “ <a href="#">Heterogeneous SAN storage system coexistence</a> ” (page 200).

**Table 126 P6000/EVA storage system rules** (continued)

Rule number	Description
6	For EVA4400 storage systems with the EVA4400 Embedded Switch Module, 8 Gb Brocade connectivity rules, see “B-series Fibre Channel switch fabric rules” (page 96).
7	<p>EVA4400 (without the embedded switch module) with XCS 09x and EVA4000/4100/6000/6100/6400/8000/8100/8400 are supported with 8 Gb/s, 4 Gb/s, or 2 Gb/s switch or HBA direct connectivity only (see rule 9).</p> <ul style="list-style-type: none"> <li>• EVA6400/8400 requires XCS 095x minimum.</li> <li>• EVA4100/6100/8100 requires XCS 6.2x minimum.</li> <li>• P6300/P6500 EVA requires XCS 10000000 minimum.</li> <li>• All Fibre Channel host port connections to P6000/EVA controllers must be set to auto, 8 Gb/s (P6300/P6500 EVA only), 4 Gb/s, or 2 Gb/s. HP recommends setting the port speed to auto.</li> <li>• Fibre Channel 1 Gb/s components or fabric segments are allowed, including 1 Gb/s links for SAN extension using long-distance direct Fibre Channel, WDM, or FCIP with 8 Gb/s, 4 Gb/s or 2 Gb/s EVA models.</li> <li>• P6000/EVA storage management server, Storage Management Appliance, and other supported server connections to P6000/EVA controllers are allowed at 1 Gb/s when connected to a Fibre Channel switch, provided that the P6000/EVA port is connected at 8 Gb/s, 4 Gb/s, or 2 Gb/s.</li> </ul> <p><b>Note:</b> This rule applies to P6000/EVA controllers connected to a Fibre Channel switch and to P6000/EVA controllers using direct connect to a server or HBA (also see rule 9).</p>
8	<p>EVA4000/4100/4400/6000/6100/6400/8000/8100/8400 supports active/active failover. EVA3000/5000 supports active/active failover (VCS 4x) or active/passive failover (VCS 3x).</p> <p>Active/active failover and active/passive failover require a minimum of two Fibre Channel HBAs and native operating system or layered multipathing driver functionality.</p> <p>For exceptions, see “Single path implementation” in the <i>HP StorageWorks 4x00/6x00/8x00 Enterprise Virtual Array User Guide</i>, available at <a href="http://www.hp.com/support/manuals">http://www.hp.com/support/manuals</a>. In the Storage section, click <b>Disk Storage Systems</b>, and then select your product.</p>
9	<p>P6000/EVA storage systems using XCS 6.x (or later) or XCS 9.x (or later) are supported in configurations with a combination of fabric attached and direct connect servers (direct connect uses no Fibre Channel switch). This is subject to the system maximum for host connections. EVA XL storage system support includes support for mixed Fibre Channel speeds on different host ports (4 Gb/s and 2 Gb/s) when using 4 Gb/s controllers.</p> <ul style="list-style-type: none"> <li>• HP-UX, Linux, and Microsoft Windows are supported for direct connect.</li> <li>• In a Linux environment, direct connect is supported with Emulex HBAs, except when booting from Itanium servers. Direct connect is not supported with Brocade or QLogic HBAs (except where listed below).</li> </ul> <p>Note the following direct connect restrictions for the EVA6400/8400:</p> <ul style="list-style-type: none"> <li>• In a Windows environment (4 and 8 Gb), direct connect is supported with Emulex HBAs, except when booting the operating system from the EVA. Direct connect and direct connect boot are supported with QLogic HBAs. Direct connect is not supported with Brocade HBAs.</li> <li>• In a Linux environment, direct connect is supported with Emulex 4 and 8 Gb HBAs, except when booting the operating system from the EVA. Direct connect and direct connect boot are supported with QLogic 8 Gb HBAs. Direct connect is not supported with Brocade HBAs.</li> <li>• In an HP-UX environment, direct connect is not supported with 2 Gb HBAs or Emulex 4 Gb PCI-e HBAs. Direct connect is not supported with Brocade HBAs.</li> </ul> <p>Use the EVA OCP or WOCP to specify the proper topology setting for the controller host ports (P6300/P6500/EVA4000/4100/6000/6100/6400/8000/8100/8400).</p> <ul style="list-style-type: none"> <li>• <b>For fabric</b>—Set the applicable HSV controller host ports to <code>Fabric</code>.</li> <li>• <b>For direct connect</b>—Set the applicable HSV controller host ports to <code>Direct Connect</code>.</li> </ul> <p>For information about the OCP or WOCP, see the P6000/EVA storage system documentation for your model.</p> <p>For information about configuring direct connect, see the <i>HP 4400 Enterprise Virtual Array Installation Guide</i>.</p>

**Table 126 P6000/EVA storage system rules** (continued)

Rule number	Description
10	All P6000/EVA host ports must contain a cable or a loopback connector; otherwise, host port error events will persist. If the P6000/EVA host port is empty, perform the following steps: <ol style="list-style-type: none"> <li>1. From the OCP or WOCP, set the port to <code>direct connect</code> mode.</li> <li>2. Insert a loopback connector when a P6000/EVA host port is not connected to a switch or an HBA (for direct connect).</li> </ol>
11	Supports connection of single HBA servers. Contact an HP storage representative for the white paper <i>Connecting Single HBA Servers to the Enterprise Virtual Array without Multipathing Software</i> . Servers without native or layered multipathing software are not supported for HP P6000 Continuous Access.
12	Overlapping zones are supported with disk and tape.
13	Overlapping storage port zones are supported if multiple operating systems share an array port.
14	Use storage system LUN presentation to enable/disable LUN access to specific hosts.
15	All host table entries must have the operating system type parameter set (based on the operating system accessing the assigned LUNs).
16	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.

For information about configuring the Storage Management Appliance with storage systems in the same SAN fabric, see “[Storage management server integration](#)” (page 221).

## Configuration parameters

This section describes general P6000/EVA configuration parameters. Specific solutions, such as high-availability clusters, or applications, such as HP P6000 Continuous Access (see “[HP P6000 Continuous Access SAN integration](#)” (page 220)), can define additional configuration parameters or requirements. Solution requirements must be observed as specified by the solution configuration documentation.

[Table 127](#) (page 217) lists the maximum number of connections supported by P6000/EVA storage, and the storage limits for each platform. The maximums are for access to a single P6000/EVA storage system with redundant dual controllers. If the connection requirements for the number of servers in a particular SAN exceed the maximum, deploy multiple storage systems in the SAN.

General maximums are as follows:

- 1,024 HBAs
- Up to 256 LUN presentations per a single virtual disk
- Up to 8,192 LUN presentations for all LUNs

For more information about P6000/EVA maximums, see the *HP P6000 Enterprise Virtual Array Compatibility Reference*.

A LUN presentation is the number of hosts presented to, regardless of how many adapters are in each host. For example, if a LUN is presented to eight hosts, that LUN has eight LUN presentations. If a LUN is presented to two hosts, that LUN has two LUN presentations.

LUNs #001 through #032 are presented to 8-node cluster	= 0256 LUN presentations
LUNs #033 through #064 are presented to 8-node cluster	= 0256 LUN presentations
LUNs #065 through #096 are presented to 8-node cluster	= 0256 LUN presentations
LUNs #097 through #128 are presented to 8-node cluster	= 0256 LUN presentations
LUNs #129 through #160 are presented to 8-node cluster	= 0256 LUN presentations
LUNs #161 through #192 are presented to 4-node cluster	= 0128 LUN presentations
LUNs #193 through #200 are presented to single host	= 0008 LUN presentations



When all LUNs are presented to all hosts, the number of LUNs multiplied by the number of hosts must not exceed 8,192.

Table 127 (page 217) lists the maximum number of EVA storage systems that can be configured on a single server. There is no limit on the maximum number of EVA storage systems in a SAN.

**Table 127 EVA single-server maximum configurations**

Storage systems	Operating systems <sup>1, 2</sup>	EVA storage systems per HBA <sup>3</sup>	LUNs per HBA target <sup>3, 4</sup>
EVA3000/5000 EVA4000/4100	NAS EFS Clustered Gateway - Linux Edition SUSE SLES 9	8	255
EVA6000/6100 EVA8000/8100	NAS EFS Clustered Gateway - Windows Edition Windows Storage Server 2003	8	64
EVA3000/5000 EVA4000/4100/4400 EVA6000/6100 EVA8000/8100 EVA6400/8400 P6300/P6500 EVA	Tru64 UNIX <sup>5</sup>	64	255
	IBM AIX	16	255
	HP-UX	32	255
	Linux (2.6 kernel)	16	255
	Linux (2.4 kernel)	16	128
	OpenVMS	32	255
	Sun Solaris	16	255
	VMware ESX 4.0, 3.x <sup>6</sup>	15	256 (per server)
	VMware ESX 2.5.x <sup>6</sup>	15	128 (per server)
	Windows 2000 <sup>7</sup> Windows 2003 Windows 2008	16	255 (per server)
	Xen <sup>8</sup>	16	255

<sup>1</sup> For the supported operating system versions, see “Heterogeneous server rules” (page 159).

<sup>2</sup> The maximum number of HBAs supported per server depends on the server model. (See the server documentation.) For Tru64 UNIX, the maximum is 64 HBAs per server when using 5.1a (or earlier), 255 HBAs per server when using 5.1b (or later).

<sup>3</sup> Multipath software may impose additional restrictions:

- Windows MPIO supports up to 32 paths per LUN.
- Secure Path for AIX supports up to 32 LUNs per HBA.

<sup>4</sup> The maximum number of configurable LUNs per HBA target may be less based on performance requirements.

<sup>5</sup> The P6300/P6500 EVA is not supported with Tru64 UNIX.

<sup>6</sup> VMware maximums are per VMware server, regardless of the number of guest operating systems. EVA6400/8400 is not supported with ESX 2.5.x. The P6300/P6500 EVA is not supported with ESX 2.5.x/3.x

<sup>7</sup> Windows 2000 is not supported with the P6300/P6500 EVA or EVA4400/6400/8400.

<sup>8</sup> EVA3000/5000 is not supported with the Xen operating system.

## P6000/EVA data migration

The P6000/EVA family of Fibre Channel storage systems supports data migration using the HP StorageWorks MPX200 Multifunction Router DMS feature. This feature provides for block (LUN) level data movement between source and destination storage systems.

## MPX200 Multifunction Router with data migration

The MPX200 Multifunction Router supports iSCSI, FCoE, data migration, and FCIP. The base functionality is iSCSI/FCoE, with the option to add either data migration or FCIP (one other license-enabled function) for standalone or concurrent operation. This section describes data migration usage and support.

- For information about FCoE, see “FCoE SAN fabrics” (page 24).
- For information about iSCSI configurations, see “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324) and “MPX200 Multifunction Router with iSCSI for XP storage” (page 330).
- For information about FCIP, see “MPX200 Multifunction Router with FCIP” (page 276).

A license is required to enable the data migration feature in the MPX200. All licenses are chassis-based, enabling data migration to be configured in both bays (slots) in a dual-blade chassis configuration.

The following licenses are available for data migration:

- **HP Storage Works MPX200 1TB Full Chassis Data Migration License**—Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 5TB Full Chassis Data Migration License**—Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 Full Chassis 1 Array Data Migration License**—Includes the license to migrate data from or to a single array using an MPX200 chassis.

---

**NOTE:** The MPX200 data migration has a unique fan-in/fan-out licensing model. Using this particular license, you can migrate data from multiple arrays to a single array or from a single array to multiple arrays.

---

Table 128 (page 219) describes the supported source and destination storage systems when using the MPX200 data migration feature. All arrays listed are supported for offline and online data migration, except where noted. Table 129 (page 219) describes the operating system support for online data migration. For information about configuring the MPX200 for data migration, see the *HP StorageWorks MPX200 Multifunction Router Data Migration Solution Guide*.

**Table 128 P6000/EVA data migration source-destination storage systems**

Source storage systems	Destination storage systems
<ul style="list-style-type: none"> <li>All HP MSA (Fibre Channel) and P6000/EVA models</li> <li>P9500/XP24000/20000, XP12000/10000</li> <li>SVSP</li> </ul>	<ul style="list-style-type: none"> <li>EVA4400/4400 with embedded switch</li> <li>EVA4000/4100/6000/6100/8000/8100</li> <li>EVA6400/8400</li> <li>P6300/P6500 EVA</li> </ul>
<p>Third-party array models:</p> <ul style="list-style-type: none"> <li>Dell EqualLogic family (iSCSI)<sup>1</sup></li> <li>EMC CLARiiON AX series, CX Series, Symmetrix families</li> <li>Hitachi Data Systems V series, AMS Family, WMS, USP families</li> <li>IBM DS3000 series, DS4000 series, DS5000 series, DS6000 series, DS8000 series</li> <li>NetApp FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<sup>1</sup></li> <li>Xiotech Emprise 5000, Mag3D 4000</li> </ul>	

<sup>1</sup> iSCSI arrays are only supported for offline data migration.

**Table 129 Online data migration operating system support**

MPX200 online data migration support <sup>1</sup>	Online data migration destination storage system and firmware (minimum)
<ul style="list-style-type: none"> <li>HP-UX 11.31, 11.23</li> <li>IBM AIX 6.1, 5.3</li> <li>Microsoft Windows 2008 and 2003, Failover Clustering and MSCS</li> <li>Red Hat 5 U4, 4 U8, U7, U6, U3</li> <li>SUSE 11, 10 SP3, SP1</li> <li>Sun Solaris 10</li> <li>VMware ESX 4.1, 4.0, 3.5, Clusters</li> </ul>	<ul style="list-style-type: none"> <li>P6500/P6300 (10000000)</li> <li>EVA8000/6000/4000 (6.200)</li> <li>EVA8100/6100/4100 (6.220)</li> <li>EVA8400/6400/4400 (09522000, 09534000) See "P6000/EVA data migration" (page 217).</li> <li>P9500 (70-00-50-00)</li> <li>XP24000/20000 (60-06-10-00)</li> <li>XP12000/10000 (50-09-83-00) See "P9000/XP data migration" (page 231)</li> <li>HP P10000 (V400/V800) (3.1.1) <b>Note:</b> Requires MPX200 FW 3.3.1.0 minimum.</li> <li>3PAR F-Class, T-Class, 2.3.1 (MU42) See "HP 3PAR data migration" (page 243).</li> </ul>

<sup>1</sup> For operating system updates, the minimum supported version is specified.

For current data migration storage system support and up-to-date operating system version support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## Data migration considerations

MPX200 connectivity to P6000/EVA storage as a data migration destination array is obtained through a Fibre Channel switch configured in the same fabric as the MPX200 Fibre Channel ports. When the data migration operation is complete, server connectivity to the P6000/EVA storage system must be configured based on current P6000/EVA support for operating systems, HBAs, and multipath software. For current operating system/version and data migration storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

# HP P6000 Continuous Access SAN integration

HP P6000 Continuous Access is supported in a heterogeneous SAN, provided that you follow the rules described in [Table 130](#) (page 220).

**Table 130 HP P6000 Continuous Access heterogeneous SAN configuration rules**

Rule number	Description
1	<p>HP strongly recommends that all HP P6000 Continuous Access deployments implement level 4 NSPOF SANs using two or more separate fabrics. See <a href="#">“Data availability”</a> (page 38).</p> <p>HP P6000 Continuous Access supports a subset of operating systems listed in this guide, which limits the types of servers that can reside in the HP P6000 Continuous Access management zone. For more information, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a>.</p>
2	<p>The maximum number of switches and switch hops supported in a fabric with HP P6000 Continuous Access is based on the limits for B-series, C-series, H-series, and M-series fabrics. All active and standby links must conform to the switch hop limits, including the host-to-local storage link, the local storage-to-remote storage link, and the local host-to-remote storage link. For more information, see the <i>HP P6000 Continuous Access Implementation Guide</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> and the switch and fabric rules. For switch and fabric rules, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">“B-series switches and fabric rules”</a> (page 83)</li> <li>• <a href="#">“C-series switches and fabric rules”</a> (page 115)</li> <li>• <a href="#">“H-series switches and fabric rules”</a> (page 131)</li> <li>• <a href="#">“M-series switches and fabric rules”</a> (page 140)</li> </ul> <p>For HP P6000 Continuous Access with XCS version 09534000 (or later), HP supports two storage system data replication protocols: the HP-FC protocol and the HP SCSI-FCP protocol. You must configure the Fibre Channel switches based on which data replication protocol you use.</p> <p>For H-series switches, you must use the HP SCSI FC Compliant Data Replication Protocol.</p>
3	<p>Shared usage of HP P6000 Continuous Access-configured storage systems by non-HP P6000 Continuous Access-configured servers (for example, a single HBA or an operating system without multipathing support) or non-HP P6000 Continuous Access-supported operating systems is not supported.</p>
4	<p>For information about supported versions of clustering software and Secure Path, contact an HP storage representative.</p>
5	<p>Each HP P6000 Continuous Access implementation can contain 16 EVAs; each EVA is limited to 256 HBAs. With 2 HBAs per server, 128 servers are possible. Multiple HP P6000 Continuous Access solutions can exist in the same SAN, provided that no solution exceeds the 16-array limit, imposed by zoning.</p>
6	<p>The maximum number of copy sets, DR groups, and remote copy sets is based on the EVA storage system model and controller software version (see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i>). On all storage systems, the limit is the total number of DR groups and copy sets that are either a source or a destination. When replicating across storage systems with different limits, the lower limit applies to the storage system replication pair.</p>
7	<p>The HP P6000 Continuous Access link supports mixed heterogeneous SAN, HP P6000 Continuous Access, and OpenVMS host-based shadowing traffic.</p>
8	<p>Two Storage Management Appliance Command View element managers are required: one active and one either active in standby mode or in powered-off, passive mode. The active appliance and HP P6000 Command View can be used for initial setup of HP P6000 Continuous Access storage. Management of the operational HP P6000 Continuous Access environment is done through the HP Replication Solutions Manager and other products installed on the Storage Management Appliances.</p> <p>For more information, see the <i>HP P6000 Continuous Access Implementation Guide</i>, available at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a>.</p>
9	<p>When adding EVA storage systems to a zone, use the controller port WWNs when implementing a straight-cable configuration. Use the storage system WWN when implementing a cross-cable configuration. For more information, see <a href="#">“Cabling”</a> (page 222) and the <i>HP P6000 Continuous Access Implementation Guide</i>.</p>

## Zoning

These zoning rules apply to a heterogeneous SAN with P6000/EVA storage:

- Zoning is required for all operating systems that access P6000/EVA storage systems. See “Heterogeneous server rules” (page 159).
- P6000/EVA storage can be in multiple operating system zones.

For configurations that support common server access to multiple storage system types, see “Common server access” (page 201).

Table 131 (page 221) lists zoning rules for heterogeneous SANs with P6000/EVA storage.

**Table 131 P6000/EVA configuration zoning rules**

Rule number	Description
1	If a storage management server resides in the fabric, each operating system type must be in a separate zone or VSAN. For more information, see “Storage management server integration” (page 221).
2	For C-series switches, each operating system type must be in a separate zone or VSAN.

## Tape storage

The following rules apply to tape storage support in a SAN with P6000/EVA storage:

- Overlapping zones are supported with disk and tape.
- Separate or common HBAs for disk and tape connections are supported.
- For a VLS, multiple heterogeneous initiators and multiple ports of the VLS device are supported in the same fabric zone.
- HP recommends using a separate tape-to-HBA connection for servers with backups that require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your HP representative.

## P6000/EVA SAN boot support

For current storage system support, see the Boot from SAN website at <http://h18006.www1.hp.com/storage/networking/bootsan.html> and the SPOCK website at <http://www.hp.com/storage/spock>. For the SPOCK website, you must sign up for an HP Passport to enable access.

For HP P6000 Continuous Access, if the operating system supports boot from SAN, replication of the boot disk is supported.

SAN boot through the B-series MP Router is not supported.

## Storage management server integration

A management server is required to manage an P6000/EVA storage system. The management server can be an SMA, GPS, management station (dedicated server), or HP Storage Server. The management server communicates with storage systems in-band through a Fibre Channel connection.

**NOTE:** Command View EVA 6.0 (or later) includes a more flexible security feature, which requires the establishment of read/write (ability to manipulate storage) and/or read-only accounts. Before using P6000 Command View or Command View EVA, be sure to review and implement the Command View account security feature setup procedures. Once you enable security, you cannot disable it.

Table 132 (page 222) describes the rules for using a management server for P6000/EVA storage systems.

**Table 132 Storage management server configuration rules**

Rule number	Description
1	A storage management server is required for any fabric that contains a P6000/EVA storage system.
2	P6000 Command View 9.4 (or later) is required for P6300/P6500 EVA.
3	Command View EVA 9.0 (or later) is required for the EVA6400/8400.
4	Command View EVA 8.0 (or later) is required for the EVA4400.
5	Command View EVA 5.0 (or later) supports management of multiple fabrics from a single instance.
6	Command View EVA 5.0 (or later) is required for the EVA iSCSI Connectivity Option.
7	Command View EVA 4x (or later) supports EVA4000/4100/6000/6100/8000/8100 storage systems using XCS 6x (or later), and EVA3000/5000 storage systems using VCS 4.100 (or later).
8	If a storage management server resides in the fabric, it is recommended that it be configured in a separate zone from all operating systems. Create a storage management zone for the storage management server and the elements it will monitor and manage.
9	It is not necessary to include the switch WWNs or server HBA WWNs in the storage management server zone. Management communication to these devices from the storage management server is done out-of-band or outside the fabric via TCP/IP.
10	EVA Element Manager can operate in a dual-fabric configuration. Up to 16 P6000/EVA storage systems can be managed from a single instance of Command View Element Manager.
11	Multiple storage management servers per fabric are allowed. A single P6000/EVA can be zoned with more than one instance of HP P6000 Command View, independent of where HP P6000 Command View is loaded (SMA, GPS, management station, or HP Storage Server).
12	Each storage system can be managed by only one active storage management server. Any standby storage management server can be powered on, but the HP P6000 Command View or HP P6000 Continuous Access user interface must not control the storage system. For more information, see the <i>HP P6000 Continuous Access Implementation Guide</i> , available at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
13	For fabrics with more than 1,024 HBAs, the HSV controller must be zoned to limit access to a maximum of 1,024 HBAs. Add zones to the fabric as needed to adhere to the 1,024 HBA limit.

## Cabling

This section describes cabling options for high-availability multipathing configurations for P6000/EVA storage systems.


### Level 4 NSPOF configuration

Figure 73 (page 223) through Figure 77 (page 225) show cabling options when implementing a level 4, high-availability, NSPOF configuration. For a description of availability levels, see “Data availability” (page 38).

Figure 73 (page 223) shows the physical connections for a cross-cable, high-availability NSPOF configuration for EVA3000/5000 VCS 3x storage systems using two separate fabrics. This cabling is the same cabling scheme used for HP P6000 Continuous Access. You can upgrade to an HP P6000 Continuous Access solution without needing to re-cable controller connections.

The cross-cable configuration shown in Figure 73 (page 223) is also supported for non-HP P6000 Continuous Access storage systems with VCS 3.x.


**Figure 73 EVA3000/5000 VCS 3x cross-cable configuration**



25129b

Figure 74 (page 223) shows the physical connections for a straight-cable, high-availability NSPOF configuration for EVA4000/4100/6000/6100 and EVA3000/5000 VCS 4x storage systems. This cabling scheme supports non-HP P6000 Continuous Access configurations and HP P6000 Continuous Access configurations with EVA4000/4100/6000/6100 and EVA3000/5000 VCS 4x storage systems. This configuration is not supported with HP P6000 Continuous Access for EVA3000/5000 VCS 3x storage systems.


**Figure 74 EVA4000/4100/6000/6100 and EVA3000/5000 VCS 4x straight-cable, high-availability configuration**



25130b

Figure 75 (page 224) shows the physical connections for a straight-cable, high-availability NSPOF configuration for EVA4400 storage systems. This cabling scheme is also supported for HP P6000 Continuous Access.


**Figure 75 EVA4400 9x straight-cable configuration**



26408a

Figure 76 (page 224) shows the cabling scheme for both non-HP P6000 Continuous Access and HP P6000 Continuous Access configurations for EVA8000/8100 storage systems.

**Figure 76 EVA8000/8100 straight-cable configuration**




25131b

Figure 77 (page 225) shows an EVA8000/8100 configuration in which all controller host ports support two independent, dual-redundant SANs. In this configuration, SAN 1 represents a dual-redundant SAN with Fabric A and Fabric B. Path failover is available between Fabric A and Fabric B. SAN 2 represents a second dual-redundant SAN with Fabric C and Fabric D. Path failover is available between Fabric C and Fabric D.

A modified version of this configuration allows for up to eight fabrics, two configured in a dual-redundant SAN with up to six independent fabrics. A minimum of two fabrics must be configured as a dual-redundant SAN to provide redundant access for the EVA management server.



**Figure 77 EVA8000/8100 two independent, dual-redundant SAN configuration**



25132a

## Dual-channel HBA configurations


Use dual-channel HBAs when the number of server PCI slots is limited. Most installations are configured as shown in [Figure 78 \(page 225\)](#) or [Figure 79 \(page 225\)](#). Both configurations are implemented using a single PCI slot to provide access to the same targets or LUNs, or to a different set of storage targets or LUNs through separate ports on the HBA.

Each dual-channel HBA provides greater performance than a single-channel HBA for a single PCI slot.

Target ranges are examples only. The number of storage controller targets and LUNs associated with each accessible target is operating system dependent.

[Figure 78 \(page 225\)](#) shows two HBA paths connected to the same Fibre Channel switch.


**Figure 78 Single PCI slot with dual-channel HBA and one switch**



25137b

The configuration shown in [Figure 79 \(page 225\)](#) provides increased availability during a single switch failure. For example, availability to a specific set of targets is increased by configuring access to targets A, B on both paths.


**Figure 79 Single PCI slot with dual-channel HBA and two switches**



25138b

Figure 80 (page 226) shows a sample NSPOF solution with two dual-channel HBAs. This availability solution is equivalent to using two single-channel HBAs. For more information, see “Data availability” (page 38).

**Figure 80 Two dual-channel HBAs (NSPOF)**



25139b

# 13 P9000/XP and VA storage system rules

This chapter describes specific rules for the following storage systems:

- P9500
- XP24000
- XP20000
- XP12000
- XP10000
- XP1024
- XP128
- XP512
- XP48
- VA7410
- VA7400
- VA7110
- VA7100

This chapter describes the following topics:

- “P9000/XP and VA storage systems” (page 227)
- “P9000/XP and VA SAN boot support” (page 230)
- “LUN Configuration and Security Manager XP support” (page 231)
- “P9000/XP data migration” (page 231)

## P9000/XP and VA storage systems

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions
- Multipath software

## Heterogeneous SAN support

P9000/XP and VA storage systems support shared access with any combination of operating systems listed in [Table 133](#) (page 227).

**Table 133 P9000/XP and VA heterogeneous SAN support**

Storage systems	Firmware version <sup>1</sup>	Switches <sup>2</sup>	Operating systems <sup>3</sup>
P9500	70x	B-series C-series H-series <sup>4</sup> M-series <sup>5</sup>	HP-UX IBM AIX Microsoft Windows OpenVMS Red Hat Linux Sun Solaris SUSE Linux VMware ESX NonStop Citrix 5.6
VA7410	HP22	B-series	HP-UX (except 4 Gb/s switches)

**Table 133 P9000/XP and VA heterogeneous SAN support** (continued)

Storage systems	Firmware version <sup>1</sup>	Switches <sup>2</sup>	Operating systems <sup>3</sup>
VA7110		C-series M-series	IBM AIX Microsoft Windows
VA7400 VA7100	A140		Red Hat Linux Sun Solaris SUSE Linux
XP20000 XP24000	60x	B-series C-series	HP-UX IBM AIX
XP12000 XP10000	50x	H-series <sup>4</sup> M-series	Microsoft Windows OpenVMS Red Hat Linux
XP1024 XP128	21x	B-series C-series	Sun Solaris SUSE Linux
XP512 XP48	01x	M-series	Tru64 UNIX VMware ESX

<sup>1</sup> Contact an HP storage representative for the latest firmware version support.

<sup>2</sup> P9500 storage systems are not supported with 2 Gb/s switches.

<sup>3</sup> For the supported operating system versions, see “Heterogeneous server rules” (page 159).

<sup>4</sup> The H-series switches have limited operating system support. See “H-series switches servers, operating systems, and storage system support” (page 135).

<sup>5</sup> Intrepid 10000/Mi10K Director support only

## Configuration rules

Table 134 (page 228) describes P9000/XP and VA storage system rules.

**Table 134 P9000/XP and VA storage system rules**

Rule number	Description
1	P9000/XP and VA storage systems are supported in all SAN fabric topology configurations described in this guide.
2	Unless otherwise specified, P9000/XP and VA storage systems can be configured in a SAN using the switch models listed in this guide. See the following tables: <ul style="list-style-type: none"> <li>• Table 53 (page 132) for H-series switches</li> <li>• Table 21 (page 85) through Table 23 (page 87) for B-series</li> <li>• Table 38 (page 118) for C-series</li> <li>• Table 61 (page 140) and Table 62 (page 141) for M-series</li> </ul> P9500 storage systems are not supported with 2 Gb/s switches. VA7xxx storage systems are not supported with 8 Gb/s or 4 Gb/s switches on HP-UX.
3	Storage ports can be accessed from heterogeneous operating system types and multiple clusters for HA and non-HA configurations. Overlapping zones are supported with disk and tape.
4	LUN Configuration and Security Manager XP and Secure Manager VA are required for LUN isolation with multiple hosts connected through a shared array port.
5	Connection to a common server is supported for P9500/XP and VA storage systems. See “Heterogeneous SAN storage system coexistence” (page 200).

**Table 134 P9000/XP and VA storage system rules** *(continued)*

Rule number	Description
6	XP24000/20000 and XP12000/10000 storage systems support F_Port, FL_Port, and NL_Port connectivity. Use the following settings based on the required connectivity type: <ul style="list-style-type: none"> <li>• Switch/Fabric (F_Port), Fabric = On, Connection = Point to Point</li> <li>• Switch/Fabric (FL_Port), Fabric = On, Connection = FC-AL</li> <li>• Direct connect (no Fibre Channel switch) (NL_Port), Fabric = Off, Connection = FC-AL</li> </ul> See “ <a href="#">HP P9000 (XP) Continuous Access configuration support</a> ” (page 300) for settings when using HP P9000 (XP) Continuous Access.
7	VA storage systems that are shared by MSCS (Windows) and MC/Service Guard (HP-UX) clusters are supported. Configuration requires assignment and securing of LUNs to the individual clusters.
8	Multiple operating systems and multiple clusters can be supported on the same fabric with appropriate zoning.
9	Host zones must contain homogeneous operating system types.
10	Overlapping storage port zones are supported if more than one operating system must share an array port.
11	Heterogeneous operating systems can share an P9000/XP array port with the appropriate host group/mode settings (see P9000/XP documentation). All others must use a dedicated P9000/XP port.
12	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.

## Zoning

These zoning rules apply to a heterogeneous SAN with P9000/XP and VA storage systems:

- Zoning is required for all operating systems that access P9000/XP and VA storage systems. See “[Heterogeneous server rules](#)” (page 159).
- P9000/XP and VA storage systems can be in multiple operating system zones.

For configurations that support common server access to multiple storage system types, see “[Common server access](#)” (page 201).

## Tape storage


The following rules apply to tape storage support in a SAN with P9000/XP and VA storage systems:

- Overlapping zones are supported with disk and tape.
- Separate or common HBAs for disk and tape connections are supported.
- For a VLS, multiple heterogeneous initiators and multiple ports of the VLS device are supported in the same fabric zone.
- HP recommends using a separate tape-to-HBA connection for servers with backups that require more than four DLT8000 tape drives or two Ultrium (LTO) tape drives.

For more information about tape storage support, contact your HP representative.

[Figure 81](#) (page 230) shows a configuration with XP and VA storage systems in the same fabric as tape storage.

**Figure 81 P9000/XP and VA storage systems with tape storage in a shared fabric**



25140b

## P9000/XP and VA SAN boot support

P9000/XP and VA LUNs can be booted from the SAN using B-series, C-series, and M-series switches. P9000/XP LUNs can be booted from the SAN using H-series switches.

SAN boot through the B-series MP Router is not supported.

Support for booting from the SAN depends on:

- PDC
- Firmware
- HBA
- Operating system and version
- Fibre Channel port speed

For support information, contact your HP representative.

Table 135 (page 231) indicates high-level boot support where at least one combination of storage system, HBA, operating system type, and switch is supported as a bootable configuration.

**Table 135 P9000/XP and VA SAN boot support by operating system<sup>1</sup>**

Switches	Storage systems	HP-UX	Linux	Windows	Tru64 UNIX <sup>2, 3</sup>	OpenVMS <sup>2</sup>	Sun Solaris	IBM AIX	IRIX	Xen
B-series C-series H-series <sup>4</sup> M-series <sup>4</sup>	P9500	•	•	•	—	•	•	•	—	—
	XP24000	•	•	•	•	•	•	•	—	•
	XP20000	•	•	•	•	•	•	•	—	•
	XP12000	•	•	•	•	•	•	•	—	•
	XP10000	•	•	•	•	•	•	•	—	•
	XP1024	•	•	•	•	•	•	•	•	—
	XP128	•	•	•	•	•	•	•	•	—
	XP512	•	•	•	•	•	•	—	•	—
	XP48	•	•	•	•	•	•	—	•	—
	VA7410	•	•	•	•	•	—	—	—	—
	VA7400	•	•	•	•	•	—	—	—	—
	VA7110	•	•	•	•	•	—	—	—	—
	VA7100	•	•	•	•	•	—	—	—	—

<sup>1</sup> **Legend:** • = supported; — = not supported

<sup>2</sup> XP24000/XP20000 boot on OpenVMS and Tru64 requires Alpha Server console 7.3 (or later).

<sup>3</sup> XP12000/10000 boot on OpenVMS and Tru64 requires Alpha Server console 6.9 (or later).

<sup>4</sup> Not all storage systems or operating systems listed are supported with H-series or M-series switches.

## LUN Configuration and Security Manager XP support

LUN Configuration and Security Manager XP:

- Is an array-based, LUN security and configuration tool
- Enables you to limit access between hosts and array LUNs
- Uses host WWNs to identify host access by LUN
- Provides LUN security at the array level to secure data
- Provides consolidated and consistent data access management, independent of the switch vendor
- Improves boot performance during `ioscan` by limiting the visibility between the host and targets

## P9000/XP data migration

The P9000/XP family of Fibre Channel storage systems support data migration using the HP StorageWorks MPX200 Multifunction Router DMS feature. This feature provides for block (LUN) level data movement between source and destination storage systems.

## MPX200 Multifunction Router with data migration

The MPX200 Multifunction Router supports iSCSI, FCoE, data migration, and FCIP. The base functionality is iSCSI/FCoE, with the option to add either data migration or FCIP (one other license-enabled function) for standalone or concurrent operation. This section describes data migration usage and support.

For information about:

- iSCSI configurations, see “MPX200 Multifunction Router with iSCSI for XP storage” (page 330) and “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324)
- FCoE, see “FCoE SAN fabrics” (page 24)
- FCIP, see “MPX200 Multifunction Router with FCIP” (page 276)

A license is required to enable the data migration feature in the MPX200. All licenses are chassis-based, enabling data migration to be configured in both bays (slots) in a dual-blade chassis configuration.

The following licenses are available for data migration:

- **HP Storage Works MPX200 1TB Full Chassis Data Migration License**—Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 5TB Full Chassis Data Migration License**—Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 Full Chassis 1 Array Data Migration License**—Includes the license to migrate data from or to a single array using an MPX200 chassis.

---

**NOTE:** The MPX200 data migration has a unique fan-in/fan-out licensing model. Using this particular license, you can migrate data from multiple arrays to a single array or from a single array to multiple arrays.

---

[Table 136 \(page 233\)](#) describes the supported source and destination storage systems when using the MPX200 data migration feature. All arrays listed are supported for offline and online data migration, except where noted. [Table 137 \(page 233\)](#) describes the operating system support for online data migration. For information about configuring the MPX200 for data migration, see the *HP StorageWorks MPX200 Multifunction Router Data Migration Solution Guide*.



**Table 136 P9500/XP data migration source-destination storage systems**

Source storage systems	Destination storage systems
<ul style="list-style-type: none"> <li>All HP MSA (Fibre Channel) and EVA models</li> <li>P9500/XP24000/20000, XP12000/10000</li> <li>SVSP</li> </ul>	<ul style="list-style-type: none"> <li>P9500/XP24000/20000/12000/10000</li> </ul>
Third-party array models: <ul style="list-style-type: none"> <li>Dell EqualLogic family (iSCSI)<sup>1</sup></li> <li>EMC CLARiiON AX series, CX Series, Symmetrix families</li> <li>Hitachi Data Systems V series, AMS Family, WMS, USP families</li> <li>IBM DS3000 series, DS4000 series, DS5000 series, DS6000 series, DS8000 series</li> <li>NetApp FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<sup>1</sup></li> <li>• XioTech Emprise 5000, Mag3D 4000</li> </ul>	

<sup>1</sup> iSCSI arrays are only supported for offline data migration.

**Table 137 Online data migration operating system support**

MPX200 online data migration support <sup>1</sup>	Online data migration destination storage system and firmware (minimum)
<ul style="list-style-type: none"> <li>HP-UX 11.31, 11.23</li> <li>IBM AIX 6.1, 5.3</li> <li>Microsoft Windows 2008 and 2003, Failover Clustering and MSCS</li> <li>Red Hat 5 U4, 4 U8, U7, U6, U3</li> <li>SUSE 11, 10 SP3, SP1</li> <li>Sun Solaris 10</li> <li>VMware ESX 4.1, 4.0, 3.5, Clusters</li> </ul>	<ul style="list-style-type: none"> <li>P6500/P6300 (10000000)</li> <li>EVA8000/6000/4000 (6.200)</li> <li>EVA8100/6100/4100 (6.220)</li> <li>EVA8400/6400/4400 (09522000, 09534000) See "P6000/EVA data migration" (page 217).</li> <li>P9500 (70-00-50-00)</li> <li>XP24000/20000 (60-06-10-00)</li> <li>XP12000/10000 (50-09-83-00) See "P9000/XP data migration" (page 231)</li> <li>HP P10000 (V400/V800) (3.1.1) <b>Note:</b> Requires MPX200 FW 3.3.1.0 minimum.</li> <li>3PAR F-Class, T-Class, 2.3.1 (MU2) See "HP 3PAR data migration" (page 243).</li> </ul>

<sup>1</sup> For operating system updates, the minimum supported version is specified.

## Data migration considerations

MPX200 connectivity to P9500/XP storage as a data migration destination array is obtained through a Fibre Channel switch configured in the same fabric as the MPX200 Fibre Channel ports. When the data migration operation is complete, server connectivity to the P9500/XP storage system must be configured based on current P9500/XP support for operating systems, HBAs, and multipath software. For current operating system/version and data migration storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

---

## 14 SVSP storage system rules

This chapter describes the HP SVSP storage system rules. It includes the following topics:

- “SVSP storage systems” (page 234)
- “SVSP supported third-party arrays” (page 238)
- “SVSP Continuous Access” (page 238)
- “Configuration parameters” (page 239)
- “SVSP best practices” (page 240)
- “SVSP SAN boot support” (page 240)
- “Storage management server integration” (page 240)

### SVSP storage systems

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions
- Multipath software

HP SVSP is a scalable platform that provides:

- Centralized management of the HP SVSP domain and homogeneous or heterogeneous storage arrays. SVSP does not support JBOD configurations.
- Multitiered data protection (to applications) enabling tiered storage
  - Data import from supported arrays
  - Thin provisioning
- Nondisruptive data provisioning
- Local and remote replication and volume management

HP SVSP improves array utilization by creating storage pools that span multiple arrays. Storage pools are managed using a single interface and are provided with a rich set of virtualization features, even though they span multiple physical devices from multiple vendors.

After a storage pool is defined, a volume of storage, called a virtual disk, can be allocated from the storage pool to a server. The virtual disk size is flexible and independent of the size of physical storage component or any logical unit in the SAN.

### Features

HP SVSP offers the following features:

- The ability to create storage pools that span up to 16 arrays
- SAN storage-based local and remote replication
- Thin provisioning for designated virtual disks, even if the array does not support thin provisioning
- A centralized VSM GUI to manage all objects in the domain, providing a single view of data on multiple arrays
- Data migration between arrays from multiple vendors
- Disaster recovery for a site, array, or virtual disk

- Bidirectional, asynchronous remote replication with automated initial normalization between source and destination (up to 150 ms one-way latency for asynchronous mirroring)

---

**NOTE:** HP SVSP uses built-in iSCSI; therefore, no additional devices are required to connect HP SVSP to the intersite IP network. For network requirements for asynchronous replication, see [Table 140 \(page 239\)](#).

---

- Synchronous mirroring across 100 km or (0.5 ms) one way over a Fibre Channel network
- Native support for up to four site campus, metropolitan, and continental disaster-tolerant configurations
- Front end virtual disk size of 1 GB to 16 TB
- Virtual disk groups to preserve write-order consistency across multiple virtual disks using the Continuous Access feature
- HP SAN API server-based CLI that allows scripted access to VSM
- I/O continuation (no application disruption) during normalization and merging, and data migration to another pool in the domain or storage array
- Multiple-path failover support
- Seamless, online, software and firmware upgrades
- Selective storage presentation for SAN-based data zoning
- Connectivity to Fibre Channel 2, 4, and 8 Gb/s switches
- Support for up to sixteen 4-port arrays, eight 8-port arrays, or some other combination up to 128 ports

For more information, see the *HP StorageWorks SAN Virtualization Services Platform QuickSpecs* and other related SVSP documentation, available for download on the HP StorageWorks SAN Virtualization Services Platform website <http://www.hp.com/go/SVSP>.

## Software

HP SVSP software consists of the following:

- HP StorageWorks SAN Virtualization Services Platform Business Copy for copies within an SVSP domain
- HP StorageWorks SAN Virtualization Services Platform Continuous Access for copies between SVSP domains
- HP StorageWorks SAN Virtualization Services Platform thin provisioning
- HP StorageWorks SAN Virtualization Services Platform Volume Manager
- HP StorageWorks Virtualization Services Manager (VSM)
- HP StorageWorks Virtualization Services Manager Command Line Interface

---

**NOTE:** For the latest product support information, including operating system support, see the product Quick Specs at <http://www.hp.com/go/svsp> or SPOCK at: <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access to SPOCK.

---

## Hardware for a single SVSP domain

Hardware requirements are as follows:

- Minimum of one EVA, MSA, or XP
- High-availability dual-fabric environment (see “[Level 4: multiple fabrics and device paths \(NSPOF\)](#)” (page 40))

- IP-based intersite links (for SVSP Continuous Access)
- Minimum of one host
- Minimum of two HBAs per host or one dual-channel HBA per host

**NOTE:** Long-distance asynchronous remote mirroring requires an additional domain at the other site and sufficient IP-based bandwidth between sites. For network requirements for asynchronous replication, see [Table 140 \(page 239\)](#).

For more information, see “Level 4: multiple fabrics and device paths (NSPOF)” (page 40).

## Heterogeneous SAN support

SVSP storage systems support shared access with any combination of operating systems listed in [Table 138 \(page 236\)](#)

**Table 138 SVSP supported operating systems**

SVSP Firmware	Switches	Operating systems
3.0.2 (or later)	B-series C-series H-series <sup>1</sup>	HP-UX IBM AIX Microsoft Windows OpenVMS Red Hat Linux Sun Solaris SUSE Linux VMware ESX  <b>Note:</b> For current storage system support, see the SPOCK website at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.

<sup>1</sup> Not all operating systems listed are supported with H-series switches.

## Configuration rules

[Table 139 \(page 236\)](#) lists configuration rules for heterogeneous SAN storage with SVSP storage systems.

**NOTE:** HP SVSP only supports native Fibre Channel. Other interfaces, such as IP, are not supported as an SVSP transport.

**Table 139 SVSP heterogeneous SAN storage rules**

Rule number	Description
1	HP requires that SVSP deployments implement level 4 NSPOF SANs using two or more separate fabrics. Each fabric must be composed of switches from a single switch vendor, compatibility mode is not supported. See “Data availability” (page 38).
2	Unless otherwise specified, SVSP storage systems can be configured in a SAN using the switch modules listed in this guide. See the following tables: <ul style="list-style-type: none"> <li>• <a href="#">Table 21 (page 85)</a> through <a href="#">Table 23 (page 87)</a> for B-series</li> <li>• <a href="#">Table 38 (page 118)</a> for C-series</li> <li>• <a href="#">Table 53 (page 132)</a> for H-series</li> </ul>
3	Fabrics must be made of a single switch series and must not be running in compatibility mode.
4	An SVSP controller must be correctly zoned in order to achieve the maximum supported configurations. See the <i>HP StorageWorks SAN Virtualization Services Platform Best Practices Guide</i> at <a href="http://www.hp.com/go/SVSP">www.hp.com/go/SVSP</a> (dependent upon the number of DPM groups).

**Table 139 SVSP heterogeneous SAN storage rules** (continued)

Rule number	Description
5	The maximum number of switches and switch hops supported in a fabric with SVSP is based on the limits for B-series, C-series, and H-series fabrics. All active and standby links must conform to the switch hop limits, including the host-to-local storage link, the local storage-to-remote storage link, and the local host-to-remote storage link. For switch and fabric rules, see: <ul style="list-style-type: none"> <li>“B-series switches and fabric rules” (page 83)</li> <li>“C-series switches and fabric rules” (page 115)</li> <li>“H-series switches and fabric rules” (page 131)</li> </ul>
6	For shared access and heterogeneous platform zoning requirements, see the <i>HP StorageWorks SAN Virtualization Services Platform Administrator Guide</i> supplied with the SVSP media.
7	SVSP is supported with 8 Gb/s, 4 Gb/s, or 2 Gb/s switch.
8	SVSP supports active/passive failover and requires a minimum of two Fibre Channel HBAs and native operating system or layered multipathing driver functionality.
9	Single initiator zones are recommended for ease of troubleshooting and configuration.
10	Overlapping storage port zones are supported if multiple operating systems share an array port.
11	All host table entries must have the operating system type parameter set (based on the operating system accessing the assigned LUNs).
12	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.

For information about configuring the Storage Management Appliance with storage systems in the same SAN fabric, see “Storage management server integration” (page 240).

## Storage rules

The following HP arrays are supported by HP SVSP:

- XP 24000/20000
- XP 12000/10000
- XP 1024/128
- XP 256 (legacy)
- XP 512/48
- EVA3000/5000
- EVA4000/4100/4400
- EVA6000/6100
- EVA8000/8100
- EVA6400/8400
- HSG80 (data import only)
- MSA1000
- MSA1500
- MSA2000
- MSA2000fc G2 (MSA2300fc)
- MSA2000fc
- P2000 G3
- 3PAR F-Class, T-Class
- P6300/P6500 EVA

---

**NOTE:** SVSP does not support LUNs presented to iSCSI hosts.

---

For more information about arrays supported with SVSP, see:

- HP SVSP documentation at <http://www.hp.com/go/svsp>
- The SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access to SPOCK.

For information about supported vendor arrays, contact an authorized HP representative.

## SVSP supported third-party arrays

SVSP supports third-party arrays (EMC, HDS, IBM, NetApps, SGI, and Sun). For more information on supported third-party arrays, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## SVSP data migration

SVSP is supported as a data migration source storage system when using the HP StorageWorks MPX200 Multifunction Router DMS feature. This feature provides for block (LUN) level data movement between source and destination storage systems.

For data migration from SVSP to P6000/EVA, see “P6000/EVA data migration” (page 217). For data migration from SVSP to P9000/XP, see “P9000/XP data migration” (page 231), and for SVSP to HP10000/3PAR, see “HP 3PAR data migration” (page 243).

## MPX200 Multifunction Router with data migration

The MPX200 Multifunction Router supports iSCSI, FCoE, data migration, and FCIP. The base functionality is iSCSI/FCoE, with the option to add either data migration or FCIP (one other license-enabled function) for standalone or concurrent operation. This section describes data migration usage and support.

- For information about FCoE, see “FCoE SAN fabrics” (page 24).
- For information about iSCSI configurations, see “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324) and “MPX200 Multifunction Router with iSCSI for XP storage” (page 330).
- For information about FCIP, see “MPX200 Multifunction Router with FCIP” (page 276).

A license is required to enable the data migration feature in the MPX200. All licenses are chassis-based, enabling data migration to be configured in both bays (slots) in a dual-blade chassis configuration.

The following licenses are available for data migration:

- **HP Storage Works MPX200 1TB Full Chassis Data Migration License**—Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 5TB Full Chassis Data Migration License**—Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 Full Chassis 1 Array Data Migration License**—Includes the license to migrate data from or to a single array using an MPX200 chassis.

---

**NOTE:** The MPX200 data migration has a unique fan-in/fan-out licensing model. Using this particular license, you can migrate data from multiple arrays to a single array or from a single array to multiple arrays.

---

For information about configuring the MPX200 for data migration, see the *HP StorageWorks MPX200 Multifunction Router Data Migration Solution Guide*.

## SVSP Continuous Access

This section describes the SVSP Continuous Access asynchronous data replication specifications and the supported minimum and maximum transmission rates when using the SVSP VSM servers. [Table 140](#) (page 239) lists the minimum network QoS requirements when using SVSP.

**Table 140 SVSP inter-site network requirements for long distance gateways**

Specification	Description
Bandwidth	Can be shared. System bandwidth must be sufficient to accommodate changes per time. Your HP representative can provide tools for modeling.
MTU of the IP network	1,500 bytes
Maximum latency <sup>1</sup>	150 ms IP network delay one-way or 300 ms round-trip
Average packet-loss ratio <sup>2</sup>	Must not exceed 0.5% averaged over a 5-minute window
Latency jitter <sup>3</sup>	Must not exceed plus or minus 10% over a 5-minute window

<sup>1</sup> Pre-existing restriction

<sup>2</sup> A high packet-loss ratio indicates the need to retransmit data across the inter-site link. Each retransmission delays transmissions queued behind the current packet, thus increasing the time to complete pending transactions.

<sup>3</sup> Latency jitter is the difference between the minimum and maximum values, and indicates how stable or predictable the network delay. The greater the jitter, the greater the variance in the delay, which lowers the performance predictability.

**NOTE:** Applications typically require more than the minimum bandwidth to meet throughput or initial full copy requirements. For example, at 10 Mb/s it could take up to two weeks to copy 1 TB of data, longer if the average available bandwidth is less than 10 Mb/s.

## Configuration parameters

This section describes general SVSP Configuration parameters. Specific solutions, such as high-availability clusters, or applications, such as Continuous Access or Business Copy can define additional configuration parameters or requirements. Solution requirements must be observed as specified by the solution configuration documentation.

General maximums are as follows:

- Number of servers at 2 HBAs per server (each 2 HBAs will be counted as an additional server)
  - 512 with one DPM group
  - 1024 with two DPM groups
  - 1536 with three DPM groups
  - 2048 with four DPM groups
- Number of initiators per FE virtual disk is 64 (32 dual-port hosts)
- Number of paths per Host per FE virtual disk is 8
- FE virtual disks:
  - 4096 with one DPM group
  - 8192 with two DPM groups
  - 12,288 with three DPM groups
  - 16,384 with four DPM groups
- BE LUN from a single BE array is 254
- BE LUN from all arrays 1,024
- FE max capacity 16 TB
- BE LUN max capacity 2 TB
- DPM FE paths:
  - 16,384 with one DPM group
  - 32,768 with two DPM groups

- 49,152 with three DPM groups
- 65,536 with four DPM groups
- DPM BE paths 4,096 maximum per DPM (regardless of pairs, all must see same storage)

---

**NOTE:** For Large LUNs, every 2 TB counts as one LUN. For example, a Large LUN of 8 TB would count as 4 LUNs against the maximum of 2,047 LUNs per EVA.

---

For additional information about SVSP maximums, see *HP StorageWorks SAN Virtualization Services Platform Release Notes* at [www.hp.com/go/SVSP](http://www.hp.com/go/SVSP).

## SVSP best practices

For SVSP best practices, see *HP StorageWorks SAN Virtualization Services Platform Best Practices Guide* at [www.hp.com/go/SVSP](http://www.hp.com/go/SVSP).

## SVSP SAN boot support

For current storage system support, see the Boot from SAN website at [http://h20272.www2.hp.com/Pages/spock2Html.aspx?htmlFile=an\\_boot\\_guides.html](http://h20272.www2.hp.com/Pages/spock2Html.aspx?htmlFile=an_boot_guides.html). You must sign up for an HP Passport to enable access.

## Storage management server integration

A management server may be required to manage a SVSP storage system. Two management interfaces are available on the SVSP 3.0.2 (or later) version, the standard using the Java-based interface and the Command View SVSP. The management server communicates with storage systems out-of-band through an Ethernet connection.

For more information about using Command View SVSP, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.



---

## 15 HP 3PAR storage system rules

This chapter describes specific rules for the following storage systems:

- HP 3PAR F200/F400
- HP 3PAR T400/T800
- HP P10000 3PAR (V400/V800)

This chapter describes the following topics:

- “HP 3PAR storage systems” (page 241)
- “HP 3PAR data migration” (page 243)
- “HP 3PAR storage management” (page 246)

### HP 3PAR storage systems

Before implementation, contact an HP storage representative for information about support for specific configurations, including the following elements:

- Storage system firmware
- HBA firmware and driver versions
- Multipath software

### Heterogeneous SAN support

HP 3PAR storage systems support shared access with any combination of operating systems listed in [Table 141 \(page 242\)](#).

**Table 141 HP 3PAR heterogeneous SAN support**

Storage systems	InForm software	Switches	Operating systems <sup>1</sup>
F200/F400 T400/T800 V400/V800	3.1.1	B-series C-series H-series <b>Note:</b> For Fibre Channel switch model and firmware support, see the <i>HP 3PAR InForm OS 3.1.1 Configuration Matrix</i> available on SPOCK at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.	HP Integrity Virtual Machines HP-UX Citrix Xen IBM AIX Microsoft Hyper-V Microsoft Windows Oracle VM Server Red Hat Linux Sun Solaris SUSE Linux VMware ESX
F200/F400 T400/T800	2.3.1	B-series C-series H-series <b>Note:</b> For Fibre Channel switch model and firmware support, see the <i>HP 3PAR InForm OS 2.3.1 Configuration Matrix</i> available on SPOCK at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.	HP Integrity Virtual Machines HP-UX Citrix Xen IBM AIX Microsoft Hyper-V Microsoft Windows Oracle VM Server Red Hat Linux Sun Solaris SUSE Linux VMware ESX

<sup>1</sup> For the supported operating system versions, see “Heterogeneous server rules” (page 159).

## Configuration rules

Table 142 (page 243) lists configuration rules for heterogeneous SAN storage with HP 3PAR storage systems.

**Table 142 HP 3PAR storage system rules**

Rule number	Description
1	HP 3PAR storage systems are supported in all Fibre Channel SAN fabric topology configurations described in this guide.
2	HP 3PAR storage systems can be configured in a SAN using the Fibre Channel switch models listed in the HP 3PAR InForm OS 2.3.1 and 3.1.1 configuration matrices available on SPOCK at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.
3	For information about shared access and heterogeneous platform zoning requirements, see “Heterogeneous SAN storage system coexistence” (page 200).
4	HBAs and multipath software for HP 3PAR storage systems are not supported for common access to other storage systems families. HP 3PAR storage systems require 1-to-1 zoning (1 initiator to 1 target). See “Zoning” (page 246).
5	Overlapping storage port zones are supported if multiple operating systems share an array port. For more information, see the 3PAR operating systems implementation guides.
6	Use storage system LUN presentation to enable/disable LUN access to specific hosts. For more information, see the 3PAR operating systems implementation guides.
7	All host table entries must have the operating system type parameter set (based on the operating system accessing the assigned LUNs). For more information, see the 3PAR operating systems implementation guides.
8	Servers that share access to the same storage LUN require special application software (such as cluster software) to preserve data.

For information about configuring HP 3PAR storage systems using the InForm Management Console, see the *HP 3PAR InForm Management Console User's Guide*.

## Configuration parameters

For information about HP 3PAR configuration parameters, including maximum server connectivity, see the HP 3PAR InForm OS 2.3.1 and 3.1.1 configuration matrices on the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

For configuration settings for the InServ ports, see the 3PAR implementation guide for each of the supported operating systems.

## HP 3PAR data migration

The HP 3PAR family of Fibre Channel storage systems supports data migration using the HP StorageWorks MPX200 Multifunction Router DMS feature. This feature provides for block (LUN) level data movement between source and destination storage systems.

### MPX200 Multifunction Router with data migration

The MPX200 Multifunction Router supports data migration for HP 3PAR storage systems as destination arrays only.

A license is required to enable the data migration feature in the MPX200. All licenses are chassis-based, enabling data migration to be configured in both bays (slots) in a dual-blade chassis configuration. The following licenses are available for data migration:

- **HP Storage Works MPX200 1TB Full Chassis Data Migration License**—Includes the license to migrate 1 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 5TB Full Chassis Data Migration License**—Includes the license to migrate 5 TB of data using an MPX200 chassis.
- **HP Storage Works MPX200 Full Chassis 1 Array Data Migration License**—Includes the license to migrate data from or to a single array using an MPX200 chassis.

---

**NOTE:** The MPX200 data migration has a unique fan-in/fan-out licensing model. Using this particular license, you can migrate data from multiple arrays to a single array or from a single array to multiple arrays.

---

Table 143 (page 245) describes the supported source and destination storage systems when using the MPX200 data migration feature. All arrays listed are supported for offline and online data migration, except where noted. Table 144 (page 245) describes the operating system support for online data migration. See the *HP StorageWorks MPX200 Multifunction Router Data Migration Solution Guide* for information about configuring the MPX200 for data migration.

**Table 143 HP 3PAR data migration source-destination storage systems**

Source storage systems	Destination storage systems
<ul style="list-style-type: none"> <li>All HP MSA (Fibre Channel) and EVA models</li> <li>P9500/XP24000/20000, XP12000/10000</li> <li>SVSP</li> </ul>	HP 3PAR F-Class, T-Class, V-Class
Third-party array models: <ul style="list-style-type: none"> <li>Dell EqualLogic family (iSCSI)<sup>1</sup></li> <li>EMC CLARiiON AX series, CX Series, Symmetrix families</li> <li>Hitachi Data Systems V series, AMS Family, WMS, USP families</li> <li>IBM DS3000 series, DS4000 series, DS5000 series, DS6000 series, DS8000 series</li> <li>NetApp FAS2000 Series, FAS3100 Series, FAS6000 Series (Fibre Channel and iSCSI)<sup>1</sup></li> <li>Xiotech Emprise 5000, Mag3D 4000</li> </ul>	

<sup>1</sup> iSCSI arrays only supported for offline data migration.

**Table 144 Online data migration operating system support**

MPX200 online data migration support <sup>1</sup>	Online data migration destination storage system and firmware (minimum)
<ul style="list-style-type: none"> <li>HP-UX 11.31, 11.23</li> <li>IBM AIX 6.1, 5.3</li> <li>Microsoft Windows 2008 and 2003, Failover Clustering and MSCS</li> <li>Red Hat 5 U4, 4 U8, U7, U6, U3</li> <li>SUSE 11, 10 SP3, SP1</li> <li>Sun Solaris 10</li> <li>VMware ESX 4.1, 4.0, 3.5, Clusters</li> </ul>	<ul style="list-style-type: none"> <li>P6500/P6300 (10000000)</li> <li>EVA8000/6000/4000 (6.200)</li> <li>EVA8100/6100/4100 (6.220)</li> <li>EVA8400/6400/4400 (09522000, 09534000) See "P6000/EVA data migration" (page 217).</li> <li>P9500 (70-00-50-00)</li> <li>XP24000/20000 (60-06-10-00)</li> <li>XP12000/10000 (50-09-83-00) See "P9000/XP data migration" (page 231).</li> <li>HP P10000 3PAR (V400/V800) (3.1.1) <b>Note:</b> Requires MPX200 FW 3.3.1.0 minimum.</li> <li>3PAR F-Class, T-Class, 2.3.1 (MU2) See "HP 3PAR data migration" (page 243).</li> </ul>

<sup>1</sup> For operating system updates, the minimum supported version is specified.

For the latest data migration storage system, operating system, and version support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

### HP 3PAR data migration considerations

MPX200 connectivity to HP 3PAR storage as a data migration destination array is obtained through a Fibre Channel switch configured in the same fabric as the MPX200 Fibre Channel ports. When the data migration is complete, server connectivity to the 3PAR storage system must be configured based on current HP 3PAR support for operating systems, HBAs, and multipath software. For more information, see the HP 3PAR operating systems implementation guides.

## Zoning

These zoning rules apply to a heterogeneous SAN with HP 3PAR storage systems:

- HP 3PAR storage systems require 1-to-1 zoning. See “Configuration rules” (page 243).
- Zoning is required for all operating systems that access HP 3PAR storage systems. See “Heterogeneous server rules” (page 159).
- HP 3PAR storage systems can be in multiple operating system zones.

## HP 3PAR SAN boot support

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access to SPOCK.

## HP 3PAR storage management

For information about HP 3PAR storage management, see the *3PAR InForm Management Console User's Guide*.

---

## 16 Enterprise Backup Solution

One of the most significant benefits of a SAN is the ability to share the SAN infrastructure for both disk and tape. With a SAN backup solution, you get all the benefits of the SAN, such as, consolidated storage, centralized management, and increased performance. Additionally, implementing a SAN backup solution lays the foundation for advanced data protection features such as serverless backup and backup to disk. The HP solution is the HP StorageWorks Enterprise Backup Solution (EBS).

The first step in deploying an EBS is to design the backup SAN configuration. Consult the HP StorageWorks Tape Tools and the EBS compatibility matrices, which are available at <http://h18000.www1.hp.com/products/storageworks/tapecompatibility.html>.

The second step is to configure your EBS using the *HP StorageWorks Enterprise Backup Solution Design Guide*, available at [http://h18000.www1.hp.com/products/storageworks/ebs/?jumpid=reg\\_R1002\\_USEN](http://h18000.www1.hp.com/products/storageworks/ebs/?jumpid=reg_R1002_USEN).

It describes the supported EBS hardware configurations and how to configure shared tape library backup in a heterogeneous SAN environment.

The third step is to install and configure your backup application. You can find rules and recommendations for the backup applications in the application guide.

For more information about EBS, see <http://www.hp.com/go/ebs>.

---

## Part IV SAN extension and bridging

SAN extension and bridging are presented in these chapters:

- “SAN extension” (page 249)
- “iSCSI storage” (page 312)
- “Network Attached Storage” (page 354)



# 17 SAN extension

SAN extension enables you to implement disaster-tolerant storage solutions over long distances and multiple sites.

This chapter describes the following topics:

- “SAN extension overview” (page 249)
- “Fibre Channel long-distance technology” (page 251)
- “Multi-protocol long-distance technology” (page 259)
- “HP multi-protocol long-distance products” (page 265)
- “HP storage replication products” (page 288)
- “Certified third-party WDM, iFCP, and SONET products” (page 311)

## SAN extension overview


A SAN extension is an extended ISL connection between switches, typically linking two sites. An extended ISL is considered to be:

- Any distance greater than:
  - 50 m for 8 Gb/s Fibre Channel
  - 150 m for 4 Gb/s Fibre Channel
  - 300 m for 2 Gb/s Fibre Channel
  - 500 m for 1 Gb/s Fibre Channel
- Any distance between a pair of WDM, FCIP, or FC-SONET products

**NOTE:** Fibre Channel distances listed here assume the use of type OM2 fiber optic cable. For the supported distances for each Fibre Channel speed and interconnect type, see “SAN fabric connectivity rules” (page 146).

You can implement a SAN extension with any Fibre Channel topology. Figure 82 (page 249) shows three SAN extension examples.

**Figure 82 SAN extension examples**



25141a

This section describes:

- “SAN extension technology” (page 250)
- “SAN-iSCSI bridging technology” (page 250)

## SAN extension technology

HP supports the following SAN extension technologies:

- “Fibre Channel long-distance technology” (page 251), including:
  - “Fiber optic transceivers” (page 251)
  - “Wavelength division multiplexing” (page 253)
  - “Extended fabric settings for Fibre Channel switches” (page 256)
- “Multi-protocol long-distance technology” (page 259), including:
  - “Fibre Channel over Internet Protocol” (page 259)
  - “Fibre Channel over SONET” (page 264)
  - “Fibre Channel over ATM” (page 265)

Table 145 (page 250) lists the SAN extension technologies and the corresponding HP SAN extension products.

**Table 145 SAN extension technologies and HP SAN extension products**

SAN extension technology	HP product
Fibre Channel using long-wave transceivers (10 km–35 km)	B-series, C-series, H-series, and M-series switches See “Long-wave transceiver distances” (page 252)
WDM (greater than 35 km to 100–500 km)	<ul style="list-style-type: none"> <li>• C-series switches with the HP CWDM solution</li> <li>• B-series, C-series, and M-series switches with certified third-party products See “Certified third-party WDM, iFCP, and SONET products” (page 311)</li> </ul>
FCIP (greater than 10 km to 20,000 km)	<ul style="list-style-type: none"> <li>• B-series 1606 SAN Extension Switch, 400 Multi-protocol Router (400 MP Router), Multi-protocol Router Blade (MP Router Blade), and Multi-protocol Router (MP Router) with B-series switches</li> <li>• C-series MDS 9506, 9509, 9513, 9216, 9216A, 9216i, and 9222i switches using the IPS-4, IPS-8, 14/2 Multiprotocol Services Module, or 18/4 Multiservice Module</li> <li>• B-series, C-series, and H-series switches with the HP MPX200 Multifunction Router with FCIP</li> <li>• B-series, C-series, and H-series switches with the HP IP Distance Gateway (mpx110)</li> <li>• Certified third-party products See “Certified third-party WDM, iFCP, and SONET products” (page 311)</li> </ul>
FC-SONET	See “Certified third-party WDM, iFCP, and SONET products” (page 311)
FC-ATM	See “Fibre Channel over ATM” (page 265)

## SAN-iSCSI bridging technology

SAN-iSCSI bridging connects Fibre Channel networks and IP networks. HP supports the following iSCSI to Fibre Channel bridging devices:

- B-series MP Router
- C-series MDS 9506, 9509, 9513, 9216, 9216A, 9216i, and 9222i embedded IP ports
- C-series IP Storage Services Modules (IPS-4, IPS-8)

- C-series MDS 14/2 Multiprotocol Services Module
- C-series MDS 18/4 Multiservice Module
- EVA iSCSI Connectivity Option
- EVA4400 iSCSI Connectivity Option
- MPX200 Multifunction Router with iSCSI
- SR2122-2

For information about bridging with iSCSI, see “iSCSI storage” (page 312).

## Fibre Channel long-distance technology

This section describes the following Fibre Channel long-distance methodology and hardware for a SAN extension:


- “Fiber optic transceivers” (page 251)
- “Wavelength division multiplexing” (page 253)
- “Extended fabric settings for Fibre Channel switches” (page 256)

### Fiber optic transceivers

Fibre Channel switches use one of the following types of fiber optic transceivers. The transceivers can be short wave or long wave. Use long-wave transceivers for SAN extension.

- The 8 Gb/s, 4 Gb/s, and 2 Gb/s transceivers are known as SFPs and use LC style connectors (Figure 83 (page 251)).


**Figure 83 LC SFP transceiver**



25142a

- The 1 Gb/s transceivers can be LC SFPs (Figure 83 (page 251)) or GBICs, which use SC style connectors (Figure 84 (page 251)).

**Figure 84 SC GBIC transceiver**



25143a

Table 146 (page 252) lists supported long-wave transceiver distances and maximum supported distances under ideal operating conditions. Some long-wave optical transceivers can transmit up to a distance of 100 km.

**Table 146 Long-wave transceiver distances**

Interconnect speed	Transceiver	Distance/rate
10 Gb/s	10 Gb/s 10 km SFP	10 km ISL at 10 Gb/s (C-series)
8 Gb/s	8 Gb/s 10 km SFP	10 km ISL at 8 Gb/s (H-series with additional buffer credits allocated) <sup>1</sup> 10 km ISL at 8 Gb/s (B-series and C-series) 3.3 km ISL at 8 Gb/s (H-series, base switch) 6.6 km ISL at 4 Gb/s (H-series, base switch) 10 km ISL at 2 Gb/s (H-series, base switch)
	8 Gb/s 25 km SFP	25 km ISL at 8 Gb/s (B-series)
4 Gb/s	4 Gb/s 10 km SFP	10 km ISL at 4 Gb/s (B-series, C-series)
	4 Gb/s 4 km SFP	4 km ISL at 4 Gb/s (C-series)
	4 Gb/s 30 km SFP (extended-reach)	30 km ISL at 4 Gb/s (B-series 8 Gb/s switches only)
2 Gb/s	2 Gb/s 35 km SFP (extended-reach)	35 km ISL at 2 Gb/s (B-series and M-series)
	2 Gb/s 10 km SFP	10 km ISL at 2 Gb/s (B-series, C-series, M-series)
1 Gb/s	1 Gb/s 10 km GBIC	10 km ISL at 1 Gb/s (B-series)
	1 Gb/s 10 km SFP	10 km ISL at 1 Gb/s (C-series, M-series)
	1 Gb/s 35 km SFP	35 km ISL at 1 Gb/s (C-series, M-series)
	1 Gb/s 100 km GBIC	100 km ISL at 1 Gb/s (B-series)

<sup>1</sup> You can use EFMS to allocate more buffer credits to ports of an H-series switch to achieve increased distance up to the limit of the SFP capability.

For detailed information about distance rules for long-wave transceivers, see:

- “Fibre Channel distance rules for 16 Gb/s and 8 Gb/s switch models (B-series, C-series, and H-series)” (page 150)
- “Fibre Channel distance rules for 4 Gb/s switch models (B-series and C-series switches)” (page 151)
- “Fibre Channel distance rules for 2 Gb/s switch models (B-series, C-series, and M-series switches)” (page 152)
- “Fibre Channel distance rules for 1 Gb/s switch models (B-series, C-series, and M-series switches)” (page 153)

**NOTE:** To ensure adequate link performance, see “Extended fabric settings for Fibre Channel switches” (page 256).

## Wavelength division multiplexing

This section describes the following:

- “WDM overview” (page 253)
- “WDM network implementation” (page 253)
- “WDM system architectures” (page 253)
- “WDM system characteristics” (page 254)
- “HP coarse wave division multiplexing” (page 254)
- “Third-party WDM products” (page 255)

### WDM overview

WDM devices extend the distance between two Fibre Channel switches. The devices are transparent to the switches and do not count as an additional hop.

To accommodate WDM devices, you must have enough Fibre Channel BB\_credits to maintain line-speed performance. WDM supports Fibre Channel speeds of 4 Gb/s, 2 Gb/s, and 1 Gb/s. When planning SAN extension, BB\_credits are an important consideration in WDM network configurations. Typical WDM implementations for storage replication include a primary and secondary path. You must have enough BB\_credits to cover the distances for both the primary path and secondary path so that performance is not affected if the primary path fails.

### WDM network implementation

WDM-based networks provide a lower-cost way to respond quickly to increased bandwidth demands and protocol changes. The quick response occurs because each wavelength is a new, full-bandwidth communications link. In many areas of the world, it is less expensive to deploy WDM devices on existing fiber than it is to install new fiber.

After implementing WDM, service providers can establish a “grow as you go” infrastructure. Service providers can expand capacity in any portion of their networks. Carriers can address areas of congestion resulting from high-capacity demands.

WDM enables you to partition and maintain dedicated wavelengths for different customers. For example, service providers can lease wavelengths (instead of an entire fiber) to their high-use business customers.

### WDM system architectures

The WDM system architectures are as follows:

- Passive (optical transmission protocol)
- Active signal amplification
- Active protocol handling

Most WDM products use one of these architectures or combine attributes of each.

Table 147 (page 254) summarizes the WDM system architectures.

**Table 147 WDM system architectures**

System architecture <sup>1</sup>	Description
Passive (optical transmission protocol)	<ul style="list-style-type: none"> <li>• Transparent to transmission protocol and data-rate independent</li> <li>• Establishes open interfaces that provide flexibility to use Fibre Channel, SONET/SDH, ATM, Frame Relay, and other protocols over the same fiber</li> <li>• Passes the optical signal without any form of signal conditioning such as amplification or attenuation</li> </ul>
Active signal amplification	<ul style="list-style-type: none"> <li>• Includes line amplifiers and attenuators that connect to other devices through fiber optic links</li> <li>• Boosts the signals that are transmitted to and received from peripheral network devices</li> <li>• Using hardware and/or software control loops, monitors power levels to ensure that the operation does not exceed the hardware's power budgets</li> </ul>
Active protocol handling	<ul style="list-style-type: none"> <li>• Offers protocol-specific capabilities for Fibre Channel, enabling digital TDM and optical multiplexing to support multiple channels on each wavelength</li> <li>• Provides network monitoring, digital retiming (to reduce timing jitter), link integrity monitoring, and distance buffering</li> <li>• May require additional and potentially costly transmission hardware when deployed in meshed networks</li> </ul> <p><b>Note:</b> HP Continuous Access products using HP FC Data Replication Protocol require in-order delivery of data replication Fibre Channel frames. Architectures, products, or protocols that do not guarantee in-order delivery are not supported.</p>

<sup>1</sup> Active protocol handling and passive protocol handling require different switch port settings. See “Port protocol setting based on the extension architecture” (page 256).

## WDM system characteristics

To help carriers realize the full potential of WDM, HP-supported WDM systems have the following characteristics:

- Use the full capacity of the existing dark fiber
- Offer component reliability, 24x7 availability, and expandability
- Provide optical signal amplification and attenuation to increase the transmitted/received signal-to-noise ratio
- Provide signal conditioning (that is, the retiming and reshaping of the optical data-carrying signal) for optimization of the bit error rate
- Offer channel add/drop capability (the ability to change the number of data channels by adding or dropping optical wavelengths on any network node)
- Allow compensation of power levels to facilitate adding or dropping channels
- Provide upgradable channel capacity and/or bit rate
- Allow interoperability through standards-compliant interfaces such as Fibre Channel, SONET, and ATM
- Convert wavelengths at each interface channel before multiplexing with other channels for transmission

## HP coarse wave division multiplexing

HP offers CWDM, which is similar to DWDM but is less expensive, less expandable (maximum of eight channels), and covers a shorter distance (up to a maximum of 100 km using the 1 Gb/s or 2 Gb/s CWDM SFP transceivers, and a maximum of 40 km using the 4 Gb/s CWDM SFP transceivers).

CWDM allows up to eight 1 Gb/s, 2 Gb/s, or 4 Gb/s channels (or wavelengths) to share a single fiber pair. Each channel uses a different color or wavelength CWDM SFP transceiver. The channels are networked using a variety of wavelength-specific multiplexers/demultiplexers or OADMs that support ring or point-to-point topologies.

The 4 Gb/s CWDM solution includes the following components:

- A 2-slot chassis for CWDM multiplexer modules
- Two 4-channel OADMs
- One 8-channel multiplexer/demultiplexer
- CWDM SFP transceivers (1470, 1490, 1510, 1530, 1550, 1570, 1590, 1610 nm)

The 1 Gb/s or 2 Gb/s CWDM solution includes the following components:

- A 2-slot chassis for OADM CWDM multiplexer modules, which have the following components:
  - Eight 1-channel OADMs
  - Two 4-channel OADMs
  - One 8-channel multiplexer/demultiplexer
- CWDM SFP transceivers (1470, 1490, 1510, 1530, 1550, 1570, 1590, 1610 nm)

A typical CWDM SFP transceiver installation includes:

- Two multiplexers/demultiplexers or two OADMs
- Up to eight matched pairs of CWDM SFP transceivers of the same frequency
- Up to eight single-mode fiber optic cables
- One long-distance, single-mode fiber optic cable

For more information, see the C-series Coarse Wave Division Multiplexer product documentation at <http://h18006.www1.hp.com/products/storageworks/cwdm/index.html>.

**NOTE:** The HP CWDM multiplexer solution is not supported on B-series, H-series, or M-series switches.


For more information about B-series, M-series, and C-series supported WDM devices, see “Certified third-party WDM, iFCP, and SONET products” (page 311).

## Third-party WDM products

HP supports third-party WDM products that have been tested successfully with a wide range of HP products. See “Certified third-party WDM, iFCP, and SONET products” (page 311).


Figure 85 (page 255) and Figure 86 (page 256) show basic WDM SAN configuration options for typical HP Continuous Access storage systems. For specific configuration options, see the HP P6000 Continuous Access and HP P9000 (XP) Continuous Access product documentation.

**Figure 85 Basic WDM configuration using one long-distance fiber optic link**



The configuration in Figure 85 (page 255) is low cost, but has no long-distance link redundancy.

**Figure 86 Fully redundant WDM configuration using two long-distance fiber optic links**



25145b

The configuration in [Figure 86 \(page 256\)](#) is high cost, but is fully redundant.

HP supports the following third-party WDM products and configurations:

- CWDM and DWDM systems supported by the switch vendors.
- WDM devices that are configurable to a data rate of 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, or 10 Gb/s
- Up to 500 km at 1 Gb/s, over a WDM link (switch model dependent). On certain Fibre Channel switches, there may be reduced performance at this distance. See [“Storage product interface, switches, and transport distance rules” \(page 150\)](#).
- Performance levels depend on the number of buffers available in the switch and the amount of application data.

For information about HP-certified third-party WDM products, see [“Certified third-party WDM, iFCP, and SONET products” \(page 311\)](#).

## Extended fabric settings for Fibre Channel switches

When extending fabrics with Fibre Channel long-distance transceivers or WDM, it is important to maintain the performance of ISL connections. For information about the B-series Extended Fabrics license, see [“Features” \(page 88\)](#).

### B-series switch settings

An Extended Fabrics license is required for B-series switches when extending a fabric beyond 10 km.

For B-series extended fabric switch settings, see the *HP StorageWorks Fabric OS Administrator's Guide*, available at <http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&lang=en&cc=us&docIndexId=179111&taskId=101&prodTypeId=12169&prodSeriesId=1143936>.

**NOTE:** The HP C-series CWDM solution is not supported on B-series switches. For B-series supported WDM devices, see [“Certified third-party WDM products” \(page 311\)](#).

[Table 148 \(page 256\)](#) describes the appropriate port protocol setting based on the WDM system architecture.

**Table 148 Port protocol setting based on the extension architecture**

WDM system architecture	B-series port protocol setting
Active protocol handling ( <a href="#">Table 147 (page 254)</a> )	portCfgISLMode slot/port, 1
Passive protocol handling	portCfgISLMode slot/port, 0



---

**NOTE:** The `portCfgISLMode` and `portCfgLongDistance` L0.5, L1, or L2 mode cannot be enabled at the same time; otherwise, fabric segmentation occurs.

The `portCfgISLMode` and `portCfgLongDistance` mode LE, LD, or LS can be enabled at the same time.

---

## B-series trunking and WDM

Consider the following when using B-series trunking with WDM:

- Trunking with WDM is supported only at 2 Gb/s or 4 Gb/s (1 Gb/s is not supported).
- Trunking de-skew timers can accommodate a maximum difference of 400 m (de-skew value of 200) between the shortest and longest ISLs in a trunk.
- Trunking distance rules are listed in the *HP StorageWorks Fabric OS Administrator Guide*.
- Trunking is not supported with ports configured for `portCfgMode slot/port, 1` (R\_RDY mode).

## C-series switch settings

Table 149 (page 257) lists the C-series switch extended fabric settings for maintaining performance with extended fabric links.

**Table 149 C-series switch extended fabric settings**

Extended fabric item	Setting
Supported switch models	C-series switches support up to 4095 BB_credits. For more information, see “C-series switches that support iSCSI” (page 347).
Maximum number of hops	7
Maximum segment distance	For more information about distance rules for long-wave transceivers, see Table 80 (page 151), Table 81 (page 152), and Table 82 (page 153) in “SAN fabric connectivity and switch interoperability rules” (page 146).

Consider the following when using C-series switches with extended fabric links:

- The MDS 9216i and 14/2 Multiprotocol Services Modules can allocate up to 3,500 BB\_credits to an individual port.
- The MDS 9222i Multiservice Fabric and 18/4 Multiservice Modules can allocate up to 4,095 BB\_credits to an individual port.
- 4 Gb/s and 10 Gb/s Fibre Channel Switching Modules can allocate up to 4,095 BB\_credits to an individual port.
- Each port on the 16-port line card supports 255 BB\_credits. These credits are available on a per-port basis through a single link or a link combined with a port channel.
- Using a port channel, bundles up to sixteen 2 Gb/s links to form a single 32 Gb/s link.
- Using a port channel, bundles up to sixteen 4 Gb/s links to form a single 64 Gb/s link.
- Using a port channel, bundles up to sixteen 10 Gb/s links to form a single 160 Gb/s link.
- All port channel links must be the same speed.

## H-series switch settings

The H-series switches have a fixed BB-credit setting. When using supported long-wave SFP, the following distances are supported:

- 3.3 km at 8 Gb/s
- 6.6 km at 4 Gb/s
- 10 km at 2 Gb/s

However, you can use EFMS to allocate more buffer credits to ports of an H-series switch to achieve increased distance up to the limit of the SFP capability, allowing 10 km at 8 Gb/s, 4Gb/s, or 2 Gb/s to be supported.

## M-series switch settings

Table 150 (page 258) lists the M-series switch extended fabric settings for maintaining performance with extended fabric links.

**Table 150 M-series switch extended fabric settings**

Extended fabric item	Setting
Supported switch models	2 Gb/s and 4 Gb/s For more information, see Table 61 (page 140) and Table 62 (page 141).
Maximum number of hops	3
Maximum segment distance	200 km between any two devices in the SAN For more information about distance rules for long-wave transceivers, see Table 80 (page 151), Table 81 (page 152), and Table 82 (page 153) in “SAN fabric connectivity and switch interoperability rules” (page 146).

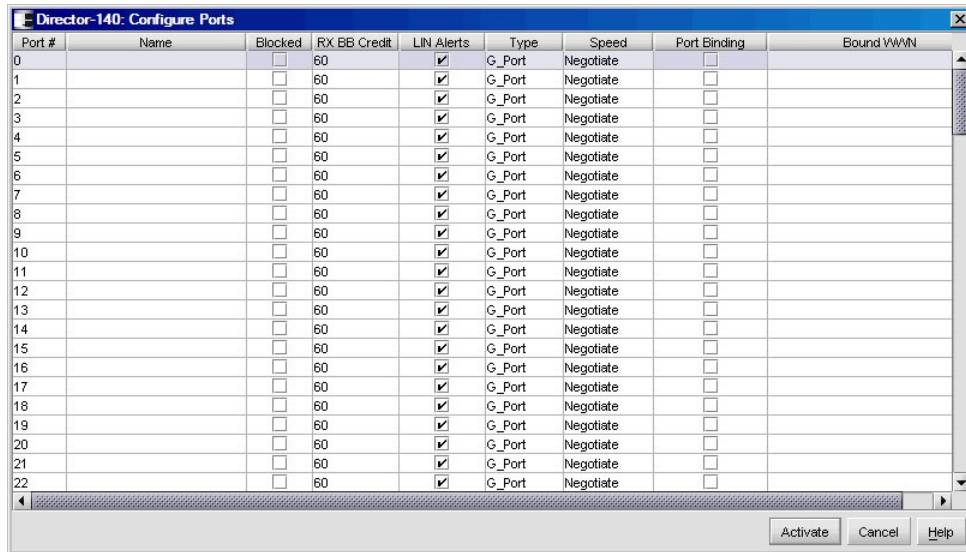
**NOTE:** The HP C-series CWDM solution is not supported on M-series switches. For M-series supported WDM devices, see “Certified third-party WDM products” (page 311).

Consider the following when using M-series switches with extended fabric links:

- The HP StorageWorks Edge Switch 2/24 has a fixed BB\_credit setting and is limited based on the port number to support links beyond 500 m for 1 Gb/s and 300 m for 2 Gb/s.
  - Ports 0 through 3 support distances up to 20 km at 1 Gb/s and up to 10 km at 2 Gb/s.
  - Ports 4 through 23 do not have enough BB\_credits to support long-wave SFP transceivers. Therefore, they only support the short-wave SFP transceiver link distances of 500 m at 1 Gb/s or 300 m at 2 Gb/s.
- To maintain line-speed performance for extended lines over 10 km, you must configure the applicable ports for the RX BB setting. The RX BB setting cannot be changed unless the port is offline.

Using HAFM, select the **Configure Ports** menu option to display the **Configure Ports** dialog box (Figure 87 (page 259)).


Figure 87 HAFM Configure Ports dialog box



Change the RX BB setting as follows:

1. Right-click in the RX BB Credit column to display the **Set All RX BB Credits** dialog box (Figure 88 (page 259)).

Figure 88 Set All RX BB Credits dialog box



2. Enter a value and click **OK** to change the values and return to the **Configure Ports** dialog box.
3. Click **Activate** to change the values on the switch.

## Multi-protocol long-distance technology

This section describes the following storage replication technologies, which enable data transfer between SAN networks:

- “Fibre Channel over Internet Protocol” (page 259)
- “Fibre Channel over SONET” (page 264)
- “Fibre Channel over ATM” (page 265)

### Fibre Channel over Internet Protocol

FCIP connects Fibre Channel fabrics over IP-based networks to form a unified SAN in a single fabric. FCIP relies on IP-based network services to provide connectivity between fabrics over LANs, MANs, or WANs.

This section describes the following topics:

- “FCIP mechanisms” (page 260)
- “FCIP link configurations” (page 260)
- “FCIP network considerations” (page 261)
- “FCIP bandwidth considerations” (page 261)

- “FCIP gateways” (page 263)
- “Third-party QoS and data encryption FCIP products” (page 263)

## FCIP mechanisms

FCIP gateways encapsulate Fibre Channel frames into IP packets and transmit them through a tunnel in an existing IP network infrastructure. The IP tunnel is a dedicated link that transmits the Fibre Channel data stream over the IP network. On the receiving end, the FCIP gateway extracts the original Fibre Channel frames from the received IP packets and then retransmits them to the destination Fibre Channel node. The gateways also handle IP-level error recovery.

**NOTE:** You must use the same gateway model (or model family in the case of MPX200 to mpx110) at both ends to ensure interoperability.

To connect to FCIP gateways, B-series and M-series switches connect through an E\_Port, while C-series switches use plug-in modules. [Figure 89 \(page 260\)](#), [Figure 90 \(page 260\)](#), and [Figure 91 \(page 261\)](#) show IP link configurations.

## FCIP link configurations

Using FCIP, you can configure the SAN ISL through a single link, dual links, or shared links.

### FCIP single-link configuration

The simplest FCIP configuration comprises one link ([Figure 89 \(page 260\)](#)).


**Figure 89 FCIP single-link configuration**



### FCIP dual-link configuration

A dual-link configuration provides redundancy ([Figure 90 \(page 260\)](#)). If one link fails, the other link temporarily handles all data replication. For enhanced fault tolerance, you can use two IP providers.

**Figure 90 FCIP dual-link configuration**




In a dual-link configuration, HP recommends that you limit the maximum sustained I/O load to 40% of the maximum available bandwidth for each link. This allows for instantaneous bursts of I/O activity and minimizes the effect of a link failure on performance.

### FCIP shared-link configuration

A shared-link configuration uses only one IP network ([Figure 91 \(page 261\)](#)).

**Figure 91 FCIP shared-link configuration**



25148b

**NOTE:** Do not use the shared-link configuration if you require high availability because it does not provide redundancy between fabrics. It can also decrease performance because the total bandwidth available for storage is shared by the two fabrics.

### FCIP network considerations

Implementing FCIP with your existing network depends on the expected storage replication application load and existing network traffic. The key consideration is whether you have enough unused or available bandwidth from your network to support the current network load, accommodate future growth, and handle replication load demands.

Table 151 (page 261) lists considerations for determining whether to use an existing network.

**Table 151 FCIP network consideration**

Configuration type	Use existing network?	Factors
Mirrored FCIP SAN	No	For peak performance, HP recommends using a separate network. A dedicated network is the benchmark for mirrored FCIP SAN systems.
Data migration	Yes	Because data migration is usually a one-time event for upgrade or maintenance purposes, you can use your existing network. However, network performance can be significantly degraded during data migration.

FCIP gateways support Ethernet connections of 10 Mb/s, 100 Mb/s, and 1 Gb/s. Select the network connection that matches the amount of data to be transferred and the time allowed for that transfer.

### FCIP bandwidth considerations

When sites are located many miles apart, there can be unacceptable delays in the completion of an I/O transaction. Increasing the available bandwidth may not solve this problem.

#### Recommendations for managing bandwidth with FCIP

In an enterprise-level SAN with multiple copy sets, merges, or full copies, normalization time can be extensive. Use the following guidelines to decrease the normalization time in an HP P6000 Continuous Access environment with the EVA family of storage arrays:

- Set up and normalize all copy sets at the same location with direct Fibre Channel connections, and then move the remote hardware to the remote site for normal operations. New copy sets will then normalize at the slower link speeds.
- Increase link bandwidth while normalization is taking place.
- Determine which data must be immediately available after a disaster, and save that data to a copy set. This applies to EVA storage arrays. Back up all other data using backup methods that can run at off-peak times.

- Most IP networks do not manage bandwidth to each connection. As traffic increases due to other demands on the network, you can use bandwidth from the replication application. Use the following techniques to minimize impact on performance:
  - Create VPNs with QoS through your local routers for the replication circuit.
  - Create separate physical networks for EVA storage arrays.
  - Guarantee the bandwidth using a third-party router and/or QoS vendor.
- Distance affects the amount of data that can be transmitted across a link. Consider these site-planning best practices:
  - Use the shortest possible distance between remote sites.
  - Always use the least possible number of copy sets. If possible, combine virtual disks that have the same failover requirements because it is best to have one copy set per application instance.
  - Copy only data that is more expensive to re-create than to copy.
  - Add copy sets that will not impact normal data traffic.
  - Consider adding controller pairs to use available bandwidth effectively.

For additional recommendations, see the *HP P6000 Continuous Access Implementation Guide*.

### **Determining the required bandwidth**

You can determine the required bandwidth for any application. This example explains how to measure the amount of new or changed data:

1. Collect the peak read and write workloads for a given period of time. For Windows operating systems, use a tool such as PERFMON to capture the current performance requirements while HP P6000 Continuous Access is not running. Similar tools exist for other operating systems.
  - At each sample interval, capture reads per second (I/Os per second), read throughput per second (Mb/s), writes per second (I/Os per second), and write throughput per second (Mb/s).
  - If possible, collect read and write latency data.
  - Perform the collection by application, capturing the data for each logical unit (device) used by that application.
2. Create a graph of each data set that shows where the peaks occur during the day.
  - Determine whether the daily average change rate is level or bursty.
  - Consider how these numbers will increase over the next 12 to 18 months. The results of this scaling process become your design goal.
  - Determine the values for RPO and RTO.
    - RPO measures how much data is lost due to a problem at the source site. By definition, an RPO of zero (no data can be lost) requires synchronous replication, regardless of which data replication product you use.
    - RTO indicates when to start using the recovery site. This measurement includes data about application failover and restart.
    - For asynchronous HP P6000 Continuous Access, the RPO design space is near zero.

- HP P9000 (XP) Continuous Access Asynchronous supports an RPO from near zero to many hours.
  - HP P6000 Continuous Access and HP P9000 (XP) Continuous Access all synchronous replication with an RTO equal to zero.
3. After the data has been collected:
- If the RPO is near zero, use the peak write rate and throughput to estimate the bandwidth you need. For some real-time applications (such as Microsoft Exchange), increase the bandwidth between 2 to 10 times this initial estimate due to wait time for link access.
  - If the RPO is greater than zero, then average the change rate over the RPO interval and use this value as an estimate of the inter-site bandwidth. You might need to increase or decrease this bandwidth, depending on the environment and the amount of time needed to complete the last write of the day before starting the next day's work.

---

**NOTE:** Because it is difficult to predict data compression before you begin, these calculations do not account for the impact of compression. If you determine that you can compress all data at a constant rate, then use that ratio to reduce the effective throughput required from the link.

---

## FCIP gateways

Table 152 (page 263) lists the HP FCIP gateways and supported switches.

**Table 152 HP FCIP gateways**

Gateway	Supported switches
<a href="#">"B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade"</a> (page 280) <a href="#">"B-series 400 MP Router and MP Router Blade"</a> (page 283) <a href="#">"B-series MP Router"</a> (page 285)	B-series switches
<a href="#">"C-series MDS 9216i, 9222i, IPS-4, IPS-8, 14/2 Multiprotocol Services Modules, 18/4 Multiservice Modules"</a> (page 287)	C-series switches
<a href="#">"HP IP Distance Gateway (mpx110)"</a> (page 266)	B-series, C-series, and H-series switches
<a href="#">"MPX200 Multifunction Router with FCIP"</a> (page 276)	B-series, C-series, and H-series switches

---

**NOTE:** For information about products supported with M-series switches, see ["Certified third-party WDM, iFCP, and SONET products"](#) (page 311).

---

## Third-party QoS and data encryption FCIP products

Third-party vendors provide two classes of FCIP solutions that complement hardware: QoS devices and IP data encryption devices. Detailed information about these products is available at the vendor websites.

### FCIP QoS products

You may need additional hardware to improve the QoS for an existing IP network. This hardware lets you use the existing network with an FCIP gateway. [Table 153 \(page 264\)](#) lists QoS hardware you can use with FCIP gateways.

**Table 153 FCIP QoS products**

Vendor name/website	Device	Purpose	Support
Allot Communications, Inc. <a href="http://www.allot.com">http://www.allot.com</a>	NetEnforcer	Application traffic and bandwidth management system	HP P6000 Continuous Access
Packeteer, Inc. <a href="http://www.packeteer.com">http://www.packeteer.com</a>	PacketShaper 6500		
Riverstone Networks, Inc. <a href="http://www.riverstonenet.com">http://www.riverstonenet.com</a>	RS 3000, RS 8000	MPLS GbE routers	

## WAN accelerator products

WAN accelerator products are not supported for use with HP Continuous Access products.

## Fibre Channel over SONET

You can connect local SANs with a SONET to create an extended SAN.

An FC-SONET gateway resides at the end of an inter-site link. Each FC-SONET gateway encapsulates Fibre Channel frames into SONET packets before transmitting the frames over the network. Upon receiving the packets, another FC-SONET gateway extracts the original Fibre Channel frames from the SONET packets and retransmits them to the destination Fibre Channel node.

The FC-SONET gateway also handles SONET-level error recovery.

This section describes the following topics:

- “FC-SONET IP link configurations” (page 264)
- “FC-SONET network considerations” (page 265)
- “Third-party SONET gateways” (page 265)


## FC-SONET IP link configurations

Using FC-SONET, you can configure the SANs through a single link, dual links, or shared ISL links.

### FC-SONET dual-link configuration

A dual-link configuration is the benchmark for disaster protection (Figure 92 (page 264)). If one link fails, the other link temporarily handles all data replication. For enhanced fault tolerance, you use two IP providers, accessing the data center through two links.

**Figure 92 FC-SONET dual-link configuration**




In a dual-link configuration, HP recommends that you limit the maximum sustained I/O load to 40% of the maximum available bandwidth for each link. This allows for instantaneous bursts of I/O activity and minimizes the effect of a link failure on performance.

### FC-SONET shared-link configuration

A shared-link configuration uses only one ISL between fabrics (Figure 93 (page 265)).



**Figure 93 FC-SONET shared-link configuration**



**NOTE:** Do not use the shared-link configuration if you require high availability because it does not provide redundancy between fabrics. It can also decrease performance because the total bandwidth available for storage is shared by the two fabrics.

## FC-SONET network considerations

Implementing FC-SONET with your existing network depends on the expected storage replication application load and existing network traffic. The key consideration is whether you have enough unused or available bandwidth from your network to support the current network load, accommodate future growth, and handle replication load demands.

HP supports the use of SONET with HP P9000 (XP) Continuous Access. See [“ATM and SONET/SDH”](#) (page 303).

## Third-party SONET gateways

For a list of HP-certified third-party SONET gateways, see [“Certified third-party WDM, iFCP, and SONET products”](#) (page 311).

## Fibre Channel over ATM

Direct FC-to-ATM conversion is supported by the Fibre Channel standards. However, currently, no vendors sell direct FC-to-ATM gateways. If you have an ATM-based network, consider using FC-to-GbE IP gateways, with an ATM blade residing on the Ethernet switch or IP router to convert the GbE to ATM.

For detailed information about distance rules for Fibre Channel over ATM, see [“ATM extension Fibre Channel distance rules”](#) (page 153).

## HP multi-protocol long-distance products

This section describes the following HP SAN extension products:

- [“HP IP Distance Gateway \(mpx110\)”](#) (page 266)
- [“MPX200 Multifunction Router with FCIP”](#) (page 276)
- [“B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade”](#) (page 280)
- [“B-series 400 MP Router and MP Router Blade”](#) (page 283)
- [“B-series MP Router”](#) (page 285)
- [“C-series MDS 9216i, 9222i, IPS-4, IPS-8, 14/2 Multiprotocol Services Modules, 18/4 Multiservice Modules”](#) (page 287)

## HP SAN extension products summary and usage

HP provides a full line of SAN extension products designed to satisfy a range of disaster recovery solutions and requirements. [Table 154](#) (page 266) describes the product features and usage for all HP supported FCIP SAN extension products. For more information, see [“Fibre Channel over Internet Protocol”](#) (page 259), [“SAN extension best practices for HP P6000 Continuous Access”](#) (page 289),

and the *Disaster Recovery Solutions: FCIP Gateway Performance Comparison* whitepaper at <http://h18006.www1.hp.com/storage/sanwhitepapers.html>.

**Table 154 Features and usage for HP supported FCIP SAN extension products**

FCIP gateway product	Supported fabrics	Inter-fabric connectivity	Supported DR software	Data compression method/use	Recommended IP bandwidths (WAN)	Network requirements
HP IP Distance Gateway (mpx110)	B-series C-series H-series	E_Port (fabric merge) EX_Port (LSAN fabric isolation) <sup>1</sup> E_Port (VSAN fabric isolation) <sup>2</sup>		Software (Use when RTT is ≥ 50 ms; or if guaranteed WAN bandwidth is ≤ 45 Mb/s)	Low: T3/DS3 (45 Mb/s) to Medium: OC-3 (155 Mb/s)	See Table 155 (page 267).
MPX200 Multifunction Router with FCIP (MPX200)	B-series C-series H-series	E_Port (fabric merge) EX_Port (LSAN fabric isolation) <sup>1</sup> E_Port (VSAN fabric isolation) <sup>2</sup>	HP P6000 Continuous Access	Software <sup>4</sup>	Low: T3/DS3 (45 Mb/s) to Medium: OC-3 (155 Mb/s) to High: OC-6 (322 Mb/s), OC-12 (622 Mb/s) up to 1 Gb/s <sup>4</sup>	See Table 156 (page 276).
B-series 1606 Extension SAN Switch DC Dir Switch MP Extension Blade	B-series	VEX_Port (LSAN fabric isolation)	HP P9000 (XP) Continuous Access <sup>3</sup>	Hardware	Low: T3/DS3 (45 Mb/s)	See Table 157 (page 281).
B-series 400 MP Router MP Router Blade	B-series	VEX_Port (LSAN fabric isolation)		Hardware	to Medium: OC-3 (155 Mb/s)	See Table 158 (page 283).
B-series MP Router	B-series	VEX_Port (LSAN fabric isolation)			to High: OC-6 (322 Mb/s), OC-12 (622 Mb/s) up to 1 Gb/s	See Table 159 (page 286).
C-series MDS 9222i, 9216i, IPS-4, IPS-8, 14/2, 18/4	C-series	E_Port (VSAN-IVR fabric isolation)		Software (IPS-4, IPS-8) Hardware (all others)		See Table 160 (page 287).

<sup>1</sup> LSAN fabric isolation is available when used with B-series switches with Integrated Routing or B-series routers with FCR.

<sup>2</sup> VSAN fabric isolation is available when used with C-series switches with inter-VSAN routing.

<sup>3</sup> HP P9000 (XP) Continuous Access is not supported by H-series switches.

<sup>4</sup> For compression usage recommendations, see the IP performance tuning section in the *MPX200 Multifunction Router User Guide*.

## HP IP Distance Gateway (mpx110)

The HP IP Distance Gateway (mpx110) provides Fibre Channel SAN extension over an IP network. Used in conjunction with the EVA or XP family of storage systems and Continuous Access software, the mpx110 enables long-distance remote replication for disaster tolerance.

Table 155 (page 267) lists the features and requirements for the IP Distance Gateway.

**Table 155 IP Distance Gateway (mpx110) features and requirements**

Feature	Requirements
Fibre Channel switches for FCIP	B-series switches—See “B-series Fibre Channel switches” (page 85) and “B-series Fibre Channel switches and routers” (page 86). C-series switches—See “C-series Fibre Channel switches for SAN-OS” (page 118). <b>Note:</b> The MDS 9020 is not supported for use with the mpx110.
IP network protocols	TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, 1,000 Mb/s Requires dedicated IP bandwidth. See Table 164 (page 291), and Table 168 (page 294).
Storage systems	HP P6000 Continuous Access HP P9000 (XP) Continuous Access
FCIP-supported operating systems	For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about using the HP IP Distance Gateway, see <a href="http://h18006.www1.hp.com/storage/networking/index.html">http://h18006.www1.hp.com/storage/networking/index.html</a> .

### IP Distance Gateway configuration examples

The IP Distance Gateway supports the configurations shown in Figure 94 (page 267) through Figure 104 (page 273).

Figure 94 (page 267) shows a basic FCIP configuration with a local mpx110 and a remote mpx110 connected through an IP WAN using one or two long-distance links.

**Figure 94 IP Distance Gateway basic FCIP configuration with one or two long-distance links**




Figure 95 (page 268) shows a high-availability configuration using pairs of mpx110s at the local site and remote site for path redundancy.

**Figure 95 IP Distance Gateway high-availability configuration with one or two long-distance links**




Figure 96 (page 268) shows a high-availability configuration that includes a redundant IP network.

**Figure 96 IP Distance Gateway high-availability configuration with a redundant IP network**




Figure 97 (page 269) shows single-path connectivity for the servers and storage. This is the lowest-cost implementation with no redundancy.

**Figure 97 IP Distance Gateway basic FCIP configuration with single-path connectivity**




Figure 98 (page 269) shows a configuration using the IP Distance Gateway and B-series switches with Integrated Routing. This provides fabric isolation between the local and remote fabrics, allowing device access without merging the fabrics. This can be implemented in all supported IP Distance Gateway configurations using B-series Fibre Channel switches with Integrated Routing or B-series routers configured for Fibre Channel routing.

**Figure 98 IP Distance Gateway FCIP with B-series Integrated Routing**




Figure 99 (page 270) shows a configuration using the mpx110 with FCIP and C-series switches with IVR. This provides fabric isolation between the local and remote fabrics, allowing device access without merging the fabrics. This can be implemented in all supported mpx110 FCIP configurations using C-series Fibre Channel switches with IVR.

**Figure 99 IP Distance Gateway FCIP with C-series IVR**




Figure 100 (page 270) is similar to Figure 95 (page 268), but offers a higher level of redundancy using additional Fibre Channel and LAN connections.

**Figure 100 Highly redundant pairs of gateways, two long distance links**

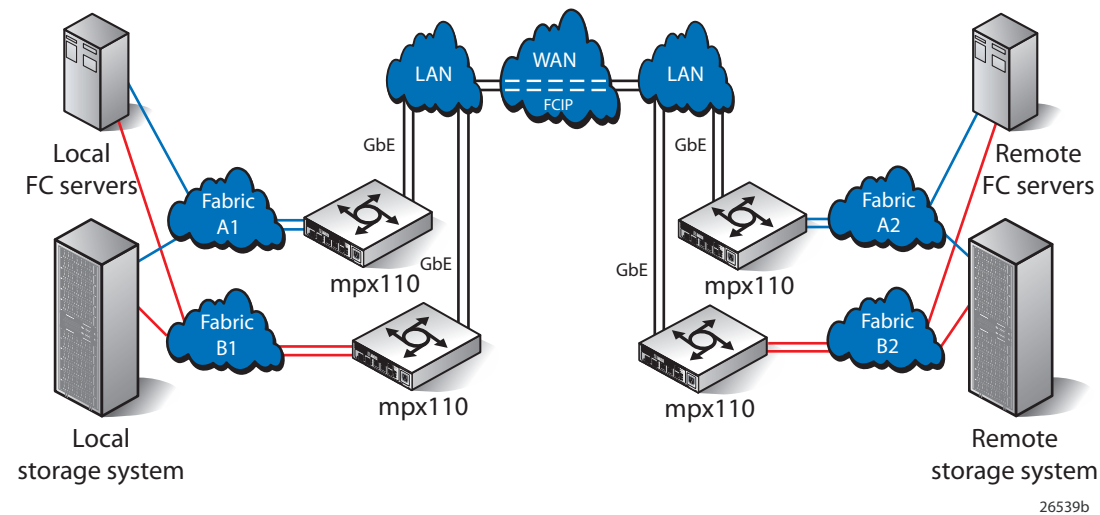



Figure 101 (page 271) is similar to Figure 96 (page 268), but offers a higher level of redundancy using additional Fibre Channel and LAN connections.

**Figure 101 Highly redundant pairs of gateways, fully redundant long-distance links**



### HP P6000 Continuous Access 3-site configurations

This section describes HP P6000 Continuous Access 3-site configurations:

- Figure 102 (page 272)
- Figure 103 (page 272)
- Figure 104 (page 273)
- Figure 105 (page 274)

The first three configurations (Figure 102 (page 272) through Figure 104 (page 273)) provide a fan-in or fan-out relationship between sites. The fourth configuration (Figure 105 (page 274)) provides for a peer-to-peer relationship between all sites.

Figure 102 (page 272) shows connectivity for three sites using four mpx110 gateways, which implements the minimum-level and lowest-cost connectivity for a 3-site configuration.

Figure 103 (page 272) shows additional connectivity and redundancy using six mpx110 gateways.

Figure 104 (page 273) shows the highest level of 3-site connectivity using eight mpx110 gateways.

The following configuration rules apply to Figure 102 (page 272) through Figure 104 (page 273) (fan-in/fan-out):

- For Site 1, Site 2 or Site 3 can function as the remote site.
- For Site 2 or Site 3, Site 1 can function as the remote site.
- Replication between Site 2 and Site 3 is not supported.


Figure 105 (page 274) is similar to Figure 103 (page 272), with additional connectivity to allow for replication between Site 2 and Site 3.

The following configuration rules apply to Figure 105 (page 274) (peer-to-peer):

- For Site 1, Site 2 or Site 3 can function as the remote site.
- For Site 2, Site 1 or Site 3 can function as the remote site.
- For Site 3, Site 1 or Site 2 can function as the remote site.

Figure 102 (page 272) shows long-distance link redundancy between all three sites.

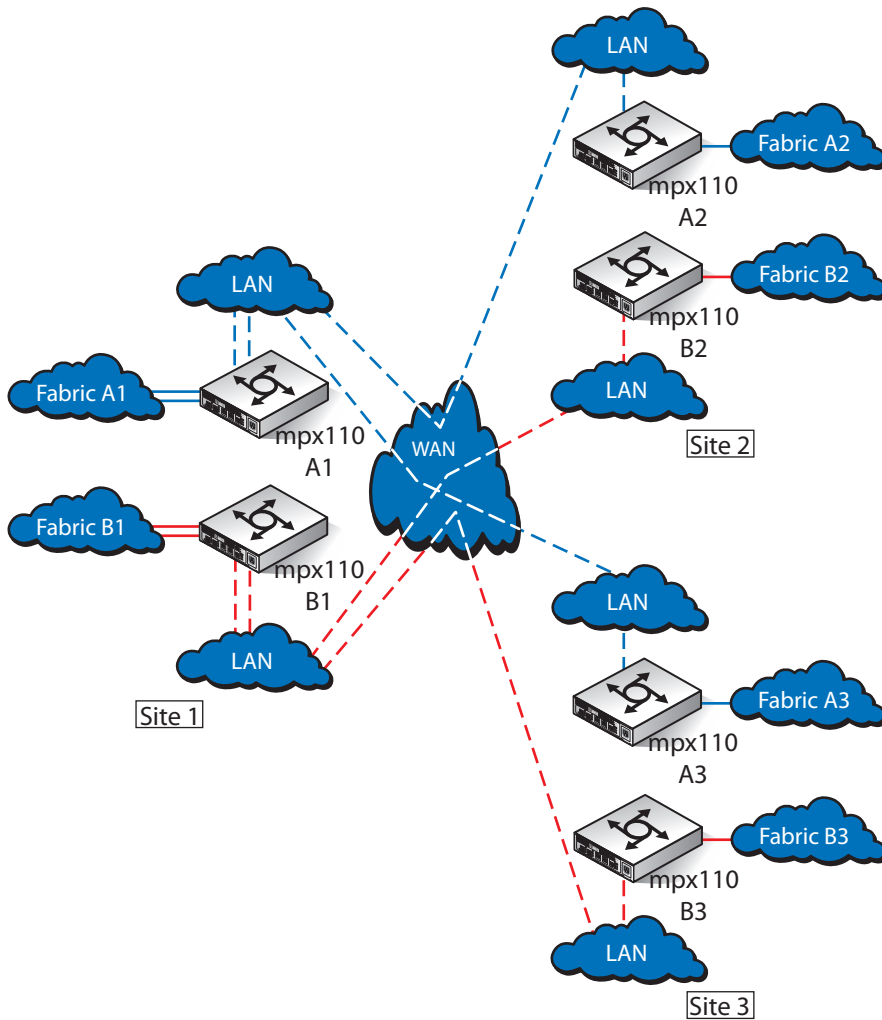
Figure 102 HP P6000 Continuous Access 3-site configuration with four mpx110 gateways



26546a

Figure 103 (page 272) shows the same long-distance link redundancy as Figure 102 (page 272), with the addition of redundant mpx110 gateways at Sites 2 and 3.

Figure 103 HP P6000 Continuous Access 3-site configuration with six mpx110 gateways




26547a



Figure 104 (page 273) shows the highest level of redundancy, with a dedicated mpx110 pair for all long-distance links to all three sites.


**Figure 104 HP P6000 Continuous Access 3-site configuration with eight gateways**



26548a

Figure 105 (page 274) shows long-distance link redundancy and full connectivity between all three sites.

**Figure 105 HP P6000 Continuous Access 3-site configuration with six mpx110 gateways, full peer-to-peer connectivity**



26582a

## Configuration rules

This section describes the configuration rules for using the mpx110 gateways for FCIP.

### General configuration rules

Review the following general configuration rules:

- All mpx110 configurations require a minimum of two mpx110 gateways, one local and one remote, connected through an IP network. These can be two mpx110s or one MPX200 with an FCIP license and one mpx110, one local and one remote, connected through an IP network. HP does not support FCIP connectivity between other gateway models.
- The mpx110 gateway is supported using FCIP extension with HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. See “EVA storage system rules” (page 275) and “HP P9000 (XP) Continuous Access” (page 296).
- Enable compression for IP fabrics with an RTT greater than or equal to 50 ms or a guaranteed WAN bandwidth of less than or equal to 45 Mb/s. For performance-tuning information based on the link speed and delay, see the *HP StorageWorks IP Distance Gateway User Guide*.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

### Operating system and multipath support

The mpx110 gateway is supported using FCIP with all operating systems and multipath software supported by HP for HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. See [Table 155 \(page 267\)](#) and [Table 156 \(page 276\)](#).

### EVA storage system rules

The EVA storage system rules follow:

- The mpx110 gateway configured for FCIP is supported for use with HP P6000 Continuous Access P6300/P6500 EVA using a minimum of XCS 10x, EVA4000/4100/6000/6100/6400/8000/8100/8400 using a minimum of XCS 6 or 09x.x, EVA4400 using a minimum of XCS 09x, and EVA3000/5000 using a minimum of VCS 4.0xx or 4.1xx.
- The mpx110 gateway is supported for use in all HP P6000 Continuous Access configurations, including the standard two-fabric, five-fabric, and six-fabric configurations. For more information, see the HP P6000 Continuous Access documentation.
- FCIP is supported for HP P6000 Continuous Access DR group LUNs and non-DR LUNs.
- Supports the minimum IP bandwidth/maximum DR groups. See [Table 166 \(page 293\)](#) through [Table 170 \(page 296\)](#).

### XP storage system rules

The XP storage system rules follow:

- The mpx110 gateway configured for FCIP is supported for use with HP P9000 (XP) Continuous Access Synchronous, Asynchronous, and Journal. For the latest support updates, see the HP storage website:  
<http://h18006.www1.hp.com/storage/saninfrastructure/index.html>
- Supported XP models are XP24000/20000 and XP12000/10000, using a minimum of 60x and 50x firmware levels respectively. Contact your HP-authorized representative for supported firmware versions.
- The mpx110 gateway is supported for use in all HP-supported HP P9000 (XP) Continuous Access FCIP configurations. For more information, see the HP P9000 (XP) Continuous Access documentation.
- Requires a minimum IP bandwidth of 16 Mb/s per path.
- For additional requirements, see “[HP P9000 \(XP\) Continuous Access FCIP gateway support](#)” ([page 307](#)).

### XP storage system software

The mpx110 gateway is supported with XP storage software applications, such as HP P9000 (XP) Continuous Access, Command View XP, HP P9000 (XP) Continuous Access Journal, Business Copy XP, and XP Array Manager.

### Fibre Channel switch and firmware support

See [Table 155 \(page 267\)](#).

## MPX200 Multifunction Router with FCIP

The HP MPX200 Multifunction Router provides Fibre Channel SAN extension over an IP network. Used in conjunction with the EVA and XP family of storage systems and Continuous Access software, the MPX200 enables long-distance remote replication for disaster tolerance.

The MPX200 FCIP feature can be configured as a standalone function or for use simultaneously with iSCSI. A license is required to enable the FCIP feature. All licenses are half-chassis based, enabling FCIP to be configured on one or both bays (slots) in a dual-blade chassis configuration.

The following licenses are available for FCIP:

- **HP Storage Works MPX200 Half Chassis FCIP License**—Includes the license to enable FCIP functionality in one of two bays (slots) in an MPX200 Chassis
- **HP Storage Works MPX200 Full Chassis FCIP License**—Includes the license to enable FCIP functionality for both bays (slots) in an MPX200 Chassis.

❗ **IMPORTANT:** If you install a single blade and a half-chassis license initially, and then install a second blade, a second half-chassis license is required.

Table 156 (page 276) lists the features and requirements for the MPX200 Multifunction Router using FCIP.

**Table 156 MPX200 Multifunction Router using FCIP features and requirements**

Feature	Requirements
Fibre Channel switches for FCIP	B-series switches—See “B-series Fibre Channel switches” (page 85) and “B-series Fibre Channel switches and routers” (page 86). C-series switches—See “C-series Fibre Channel switches for SAN-OS” (page 118). <b>Note:</b> The MDS 9020 is not supported for use with the MPX200.
IP network protocols	TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, 1,000 Mb/s Requires dedicated IP bandwidth. See Table 164 (page 291) and Table 168 (page 294).
Storage systems	HP P6000 Continuous Access HP P9000 (XP) Continuous Access
FCIP-supported operating systems	All operating systems supported for HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about using the MPX200 Multifunction Router, see <a href="http://h18006.www1.hp.com/storage/networking/index.html">http://h18006.www1.hp.com/storage/networking/index.html</a> .

## MPX200 Multifunction Router FCIP configuration examples

The MPX200 Multifunction Router supports iSCSI, data migration, and FCIP. The base functionality is iSCSI, with the option to add either data migration or FCIP (one other license-enabled function) for standalone or concurrent operation, data migration or FCIP. All FCIP configurations shown are supported with FCIP only or with simultaneous FCIP and iSCSI operations. For information about iSCSI configurations, see “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324). For information about data migration, see “MPX200 Multifunction Router with data migration” (page 218).

**NOTE:** The MPX200 does not support simultaneous FCIP and data migration operation.

The MPX200 Multifunction Router supports the FCIP configurations shown in Figure 106 (page 277), Figure 109 (page 278), and Figure 110 (page 279), in addition to all configurations shown for

the IP Distance Gateway. See “IP Distance Gateway configuration examples” (page 267). Figure 110 (page 279) shows an FCIP configuration using an MPX200 at the local site and an IP Distance Gateway at the remote site.

Figure 106 (page 277) shows a basic FCIP configuration with a local single-blade MPX200 chassis and a remote single-blade MPX200 chassis connected through an IP WAN using one or two long-distance links.

**Figure 106 MPX200 basic FCIP configuration with one or two long-distance links**




Figure 107 (page 277) shows a configuration using the MPX200 with FCIP and B-series switches with Integrated Routing. This provides fabric isolation between the local and remote fabrics, allowing device access without merging the fabrics. This can be implemented in all supported MPX200 FCIP configurations using B-series Fibre Channel switches with Integrated Routing or B-series routers configured for Fibre Channel routing.

**Figure 107 MPX200 FCIP with B-series Integrated Routing**




Figure 108 (page 278) shows a configuration using the MPX200 with FCIP and C-series switches with IVR. This provides fabric isolation between the local and remote fabrics, allowing device access without merging the fabrics. This can be implemented in all supported MPX200 FCIP configurations using C-series Fibre Channel switches with IVR.

**Figure 108 MPX200 FCIP with C-series IVR**

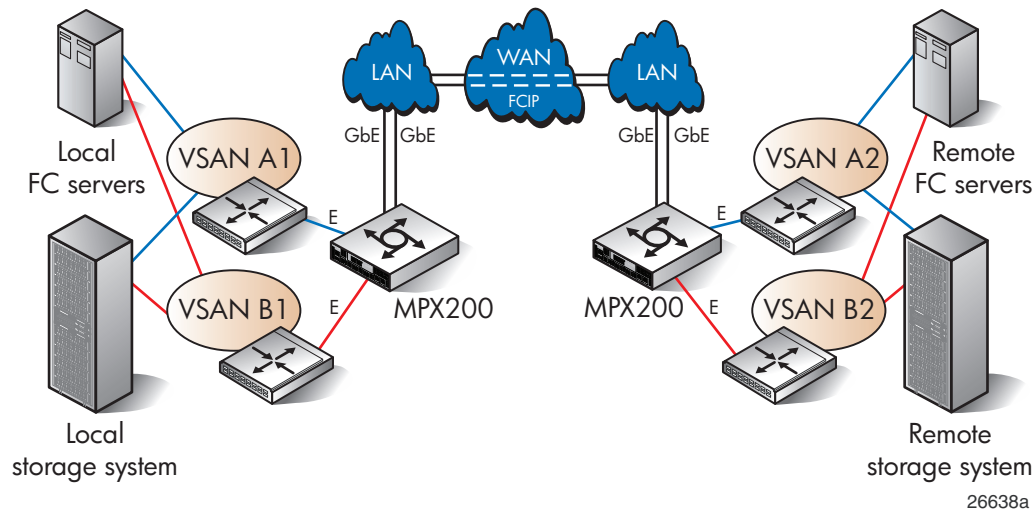


Figure 109 (page 278) shows a high-availability configuration using a dual-blade MPX200 chassis at the local site and the remote site for hardware and path redundancy.

**Figure 109 MPX200 high-availability configuration with one or two long-distance links**





Figure 110 (page 279) shows a basic FCIP configuration with a local single-blade MPX200 chassis and a remote IP Distance Gateway (mpx100).

**Figure 110 Local MPX200 basic FCIP configuration with remote IP Distance Gateway (mpx100)**



## FCIP Configuration rules

The section describes the FCIP configuration rules for using the MPX200 for FCIP.

### General FCIP configuration rules

Observe the following general configuration rules:

- All MPX200 FCIP configurations require a minimum of two gateways. These can be two MPX200s or one MPX200 and one IP Distance Gateway (mpx110), one local and one remote, connected through an IP network. HP does not support FCIP connectivity between other gateway models.
- FCIP is supported on GbE ports only (GE1 and GE2).
- The MPX200 is supported using FCIP extension with HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. See “EVA storage system rules” (page 275) and “XP storage system rules” (page 275).
- For performance-tuning information based on the link speed and delay, see the *HP StorageWorks Multifunction Router User Guide*.

For current storage system support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

### Operating system and multipath support

The MPX200 Multifunction Router supports FCIP with all operating systems and multipath software supported by HP for HP P6000 Continuous Access and HP P9000 (XP) Continuous Access. See Table 155 (page 267) and Table 156 (page 276).

### Storage system rules

Observe the following storage system rules:

- The MPX200 Multifunction Router configured for FCIP is supported for use with HP P6000 Continuous Access and P6300/P6500 EVA using a minimum of XCS 10x, EVA4000/4100/6000/6100/6400/8000/8100/8400 using a minimum of XCS 6 or

09x.x, and EVA4400 using a minimum of XCS 09x and EVA3000/5000 using a minimum of VCS 4.0xx or 4.1xx.

For XP, the MPX200 is supported for use with HP P9000 (XP) Continuous Access for XP24000/20000 and XP12000/10000.

- The MPX200 Multifunction Router is supported for use in all supported HP P6000 Continuous Access SAN configurations, including the standard two-fabric, five-fabric, and six-fabric configurations. For more information, see the HP P6000 Continuous Access documentation.
- FCIP is supported for HP P6000 Continuous Access DR group LUNs and non-DR LUNs.
- The minimum IP bandwidth required for HP P9000 (XP) Continuous Access is 16 Mb/s per path.
- Supports the minimum IP bandwidth/maximum EVA DR groups. See [Table 166 \(page 293\)](#) through [Table 170 \(page 296\)](#).

## B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade

The B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade offer FCIP SAN extension, Fibre Channel routing, and Fibre Channel switching. You can use all of these functions on the same 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade simultaneously.

A 1606 Extension SAN Switch can have:

- Up to 6 GbE IP ports for FCIP  
See [“SAN extension” \(page 249\)](#)
- Up to 16 F\_Ports for FC switching
- Up to 16 EX\_Ports for Fibre Channel routing services  
See [“SAN fabric topologies” \(page 23\)](#)

A DC Dir Switch MP Extension Blade can have:

- Up to 10 1-GbE IP ports for FCIP  
See [“SAN extension” \(page 249\)](#)
- Up to 2 10-GbE IP ports for FCIP  
See [“SAN extension” \(page 249\)](#)
- Up to 12 EX\_Ports for Fibre Channel routing services  
See [“SAN fabric topologies” \(page 23\)](#)
- Up to 12 F\_Ports for Fibre Channel switching

For the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade, you can configure a Fibre Channel port as an F\_Port, FL\_Port, E\_Port, or EX\_Port (Fibre Channel routing), and configure a GbE port as a VE\_Port (FCIP) or VEX\_Port (FCIP with Fibre Channel routing). Traffic switching is supported between the following port types:

- F\_Port, FL\_Port, E\_Port (Fibre Channel switching)
- F\_Port, FL\_Port, EX\_Port (Fibre Channel switching and Fibre Channel routing)
- F\_Port, FL\_Port, VE\_Port, VEX\_Port (Fibre Channel switching and FCIP)
- VE\_Port, VEX\_Port (Fibre Channel routing and FCIP)

---

**NOTE:** VEX\_Ports and EX\_Ports cannot connect to the same edge fabric.

VEX\_Ports are not supported with the DC Dir Switch MP Extension Blade.

---

Using FCIP and Fibre Channel routing, you can connect to local and remote fabrics without fully merging them. This prevents unauthorized access to all devices on the local and remote fabrics.



## 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements

Table 157 (page 281) lists the features and requirements for the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade.


**Table 157 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade features and requirements**

Feature	Requirements
Fibre Channel switches for FCIP	The 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade are supported with the B-series switches listed in “B-series switches and fabric rules” (page 83). These products are also supported for use as Fibre Channel switches. This provides both Fibre Channel switch connectivity for hosts and storage systems and FCIP connectivity for SAN extension.
IP network protocols	TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, and 1,000 Mb/s All 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade FCIP implementations must have dedicated IP bandwidth configured as a committed bandwidth, or using the Adaptive Rate Limiting feature. HP does not support the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade configurations with FCIP tunnels configured with nondedicated bandwidth. For more information about configuring and monitoring FCIP extension services, see <i>HP StorageWorks Fabric OS 6.x Administrator's Guide</i> . For more information on dedicated IP bandwidth requirements, see Table 164 (page 291) and Table 168 (page 294). <b>Notes:</b> <ul style="list-style-type: none"> <li>HP P9000 (XP) Continuous Access storage systems require a minimum of 16 Mb/s IP bandwidth.</li> <li>FCIP is supported with IPsec data encryption (license required). This requires a minimum of firmware 6.3.0.</li> <li>FCIP is supported with FCIP FastWrite acceleration.</li> <li>The connection to the 1606 Extension SAN Switch or DC Dir Switch MP Extension Blade FCIP port must be GbE compatible (10-GbE for the 10-GbE ports).</li> </ul>
Storage systems	HP P6000 Continuous Access HP P9000 (XP) Continuous Access (requires a minimum of firmware 6.3.0a) For operating system support, see “Heterogeneous server rules” (page 159). FCIP FastWrite is supported with HP P9000 (XP) Continuous Access and is not supported with HP P6000 Continuous Access. Supported routing modes: <ul style="list-style-type: none"> <li>HP P9000 (XP) Continuous Access—Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3). See Table 183 (page 307).</li> <li>HP P6000 Continuous Access—Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later).</li> </ul>
FCIP-supported operating systems	For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about the 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade, see the SAN Infrastructure website: <a href="http://h18006.www1.hp.com/storage/saninfrastructure/index.html">http://h18006.www1.hp.com/storage/saninfrastructure/index.html</a> .

## 1606 Extension SAN Switch configuration examples


When connecting fabrics through IP, the 1606 Extension SAN Switch serves as an FCIP gateway with Fibre Channel routing. 1606 Extension SAN Switches that communicate over FCIP links can be installed in multiple pairs for high availability (Figure 111 (page 282)) or as a single pair (Figure 112 (page 282)).

**Figure 111 NSPOF configuration with 1606 Extension SAN Switches providing Fibre Channel routing and FCIP**



26583b

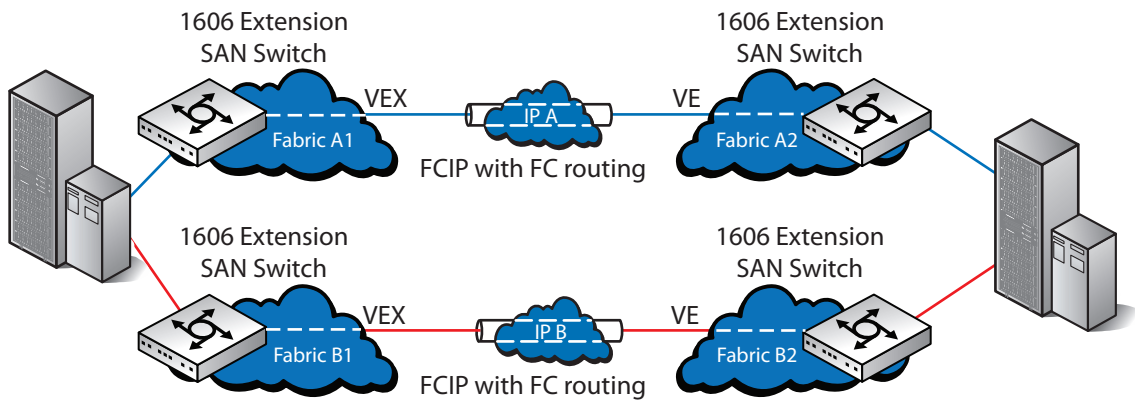
**Figure 112 Fibre Channel routing and FCIP using two 1606 Extension SAN Switches**



26584a

HP supports FCIP configurations in which the 1606 Extension SAN Switch serves as an FCIP gateway and a Fibre Channel switch. Servers and storage systems that support Continuous Access with FCIP can be directly connected to the Fibre Channel ports on the 1606 Extension SAN Switch (Figure 113 (page 282)).

**Figure 113 NSPOF configuration with 1606 Extension SAN Switches providing Fibre Channel routing and FCIP with direct connect devices**



26585b

## B-series 400 MP Router and MP Router Blade

The B-series 400 MP Router and MP Router Blade offer FCIP SAN extension, Fibre Channel routing, and Fibre Channel switching. You can use all three of these functions on the same 400 MP Router or MP Router Blade simultaneously. A 400 MP Router or MP Router Blade can have:

- Up to 2 GbE IP ports for FCIP. For more information, see “SAN extension” (page 249).
- Up to 16 EX\_Ports for Fibre Channel routing services. For more information, see “SAN fabric topologies” (page 23).
- Up to 16 F\_Ports for Fibre Channel switching.

For the 400 MP Router and MP Router Blade, you can configure a Fibre Channel port as an F\_Port, FL\_Port, E\_Port, or EX\_Port (Fibre Channel routing), and configure a GbE port as a VE\_Port (FCIP) or VEX\_Port (FCIP with Fibre Channel routing). Traffic switching is supported between the following port types:

- Fibre Channel switching (F\_Port, FL\_Port, E\_Port)
- Fibre Channel switching and Fibre Channel routing (F\_Port, FL\_Port, EX\_Port)
- Fibre Channel switching and FCIP (F\_Port, FL\_Port, VE\_Port, VEX\_Port)
- Fibre Channel routing and FCIP (VE\_Port, VEX\_Port)

---

**NOTE:** VEX\_Ports and EX\_Ports cannot connect to the same edge fabric.

---

Using FCIP and Fibre Channel routing, you can connect to local and remote fabrics without fully merging them. This prevents unauthorized access to all devices on the local and remote fabrics.

## 400 MP Router and MP Router Blade features and requirements

Table 158 (page 283) lists the features and requirements for the 400 MP Router and MP Router Blade.

**Table 158 400 MP Router and MP Router Blade features and requirements**

Feature	Requirements
Fibre Channel switches for FCIP	<p>The 400 MP Router and MP Router Blade are supported with the B-series switches listed in “B-series switches and fabric rules” (page 83).</p> <p>These products are also supported for use as Fibre Channel switches. This provides both Fibre Channel switch connectivity for hosts and storage systems and FCIP connectivity for SAN extension.</p>
IP network protocols	<p>TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, and 1,000 Mb/s</p> <p>All 400 MP Router and MP Router Blade FCIP implementations must have dedicated IP bandwidth. HP does not support 400 MP Router or MP Router Blade configurations with FCIP tunnels configured with nondedicated bandwidth. For more information, see “Configuring and monitoring FCIP extension services” in the <i>HP StorageWorks Fabric OS 6.x Administrator Guide</i>.</p> <p>For more information on dedicated IP bandwidth requirements, see Table 164 (page 291) and Table 168 (page 294).</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• HP P9000 (XP) Continuous Access storage systems require a minimum of 16 Mb/s IP bandwidth.</li> <li>• FCIP is supported with IPsec data encryption (license required) or FCIP FastWrite acceleration. This requires a minimum of firmware 5.2.0a. For firmware version 6.3.x, version 6.3.0c (or later) is required for FCIP FastWrite.</li> <li>• IPsec and FCIP FastWrite are mutually exclusive and cannot be configured simultaneously.</li> <li>• The connection to the 400 MP Router and MP Router Blade FCIP port must be GbE compatible.</li> </ul>

**Table 158 400 MP Router and MP Router Blade features and requirements** (continued)


Feature	Requirements
Storage systems	<ul style="list-style-type: none"> <li>• HP P6000 Continuous Access</li> <li>• HP P9000 (XP) Continuous Access</li> </ul> For operating system support, see “Heterogeneous server rules” (page 159). FCIP FastWrite is supported with HP P9000 (XP) Continuous Access and is not supported with HP P6000 Continuous Access. Supported routing modes: <ul style="list-style-type: none"> <li>• HP P9000 (XP) Continuous Access—Supported with port-based routing (aptpolicy = 1) or exchange-based routing (aptpolicy = 3). See Table 183 (page 307).</li> <li>• HP P6000 Continuous Access—Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later).</li> </ul>
FCIP-supported operating systems	For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about the 400 MP Router and MP Router Blade, see the SAN Infrastructure website: <a href="http://h18006.www1.hp.com/storage/saninfrastructure/index.html">http://h18006.www1.hp.com/storage/saninfrastructure/index.html</a> .

### 400 MP Router and MP Router Blade configuration examples

When connecting fabrics through IP, the 400 MP Router or MP Router Blade serve as FCIP gateways with Fibre Channel routing. Routers that communicate over FCIP links can be installed in multiple pairs for high availability (Figure 114 (page 284)) or as a single pair (Figure 115 (page 284)).


**NOTE:** Single-pair configurations require 400 MP Router or MP Router Blade firmware 6.1.0a (or later).

**Figure 114 NSPOF configuration with 400 MP Routers providing Fibre Channel routing and FCIP**



25270b


**Figure 115 Fibre Channel routing and FCIP using two 400 MP Routers**



26469b

HP supports FCIP configurations in which the 400 MP Router or MP Router Blade serves as an FCIP gateway and a Fibre Channel switch. Servers and storage systems that support Continuous Access with FCIP can be directly connected to the Fibre Channel ports on the 400 MP Router or MP Router Blade (Figure 116 (page 285)).

**Figure 116 NSPOF configuration with 400 MP Routers providing Fibre Channel routing and FCIP with direct connect devices**



For use-case configuration examples, see the 400 MP Router and MP Router Blade use-case white papers at [http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/storage/sanwhitepapers.html?jumpid=reg_R1002_USEN).

## B-series MP Router

The B-series MP Router offers FCIP SAN extension, Fibre Channel routing, and iSCSI bridging functionality. You can use all three capabilities on the same MP router simultaneously. An MP Router can have:

- Up to 12 GbE IP ports for either FCIP or iSCSI bridging  
See “Fibre Channel routing” (page 43)
- Up to 16 EX\_Ports for Fibre Channel routing  
See “SAN fabric topologies” (page 23)

For the MP Router, you can configure a port as a Fibre Channel E\_Port, EX\_Port (Fibre Channel routing), GbE iSCSI, or GbE FCIP. Traffic switching is supported between the following port types:

- Fibre Channel and Fibre Channel routing (E\_Port, EX\_Port)
- Fibre Channel and iSCSI (E\_Port, F\_Port, GbE iSCSI)
- Fibre Channel and FCIP (E\_Port, EX\_Port, GbE FCIP)
- Fibre Channel routing and FCIP (EX\_Port, GbE FCIP)

Traffic switching between other port combinations is not supported. This applies to port combinations on the same platform or independent platforms. For example, traffic from an iSCSI initiator connected to an MP Router cannot be routed to an EX\_Port on the same router or independent router.

Combining FCIP with Fibre Channel routing provides an FCIP configuration that lets you connect to local and remote fabrics without fully merging them. This prevents unauthorized access to all devices on the local and remote fabrics.

**NOTE:** You can use FCIP and iSCSI on the same MP Router but not on the same GbE port.

## B-series MP Router features and requirements

Table 159 (page 286) lists the features and requirements for the MP Router.

**Table 159 B-series MP Router features and requirements**


Feature	Requirements
Switches for iSCSI and FCIP	The MP Router is supported with the B-series switches listed in “B-series switches and fabric rules” (page 83). The MP Router is not supported for use as a Fibre Channel switch. Direct device connectivity to MP Router ports (F_Ports) is supported only with iSCSI See “iSCSI storage” (page 312).
IP network protocols	TCP/IP IPv4, Ethernet 10 Mb/s, 100 Mb/s, and 1,000 Mb/s Requires dedicated IP bandwidth, see Table 164 (page 291) and Table 168 (page 294). <b>Note:</b> You can only use optical GbE compatible connections to the FCIP or iSCSI ports.
Storage systems	<ul style="list-style-type: none"> <li>HP P6000 Continuous Access</li> <li>HP P9000 (XP) Continuous Access</li> </ul> For operating system support, see “Heterogeneous server rules” (page 159). Supported routing modes: <ul style="list-style-type: none"> <li>HP P9000 (XP) Continuous Access—Supported with port-based routing (<code>portcfgfcip = 2</code>) or exchange-based routing (<code>portcfgfcip = 1</code>). See Table 183 (page 307).</li> <li>HP P6000 Continuous Access—Supported with port-based routing (all XCS versions) or exchange-based routing (XCS 09534000 or later).</li> </ul> <b>Note:</b> Use the <code>portCfgFcip</code> command to set the routing mode. FCIP <code>portcfgfcip [-m LOAD_BAL (1-exchange, 2-SID_DID, 3-none)]</code>
FCIP-supported operating systems	For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about the MP Router, see the MP Router website: <a href="http://h18006.www1.hp.com/products/storageworks/mprouter/index.html">http://h18006.www1.hp.com/products/storageworks/mprouter/index.html</a> .

**B-series MP Router configuration examples**

This section describes MP Router sample configurations.

Figure 117 (page 286) shows how MP Routers are configured for FCIP and Fibre Channel routing.

**Figure 117 MP Routers connecting fabrics in an IP network**




Note the following:

- All MP Router FCIP configurations require a minimum level of IP bandwidth for HP P9000 (XP) Continuous Access and HP P6000 Continuous Access replication (see “HP storage replication products” (page 288)).
- HP P6000 Continuous Access supports multiple links, but not those that are shared using dynamic load balancing or exchange-based routing.

Figure 118 (page 287) shows Fibre Channel fabrics connecting through two IP subnets. Each MP Router uses a Fibre Channel interface to connect to the fabrics and a GbE interface to connect to the IP subnets.


**Figure 118 MP Routers connecting fabrics using two IP subnets**



In this sample configuration, MP Router 1 must have a GbE connection to IP subnet A and a GbE connection to a IP subnet B. A Fibre Channel host or device connected to Fabric 1 can access devices in Fabric 2 and/or Fabric 3.

Figure 119 (page 287) shows a redundant FCIP configuration in which pairs of MP Routers provide full redundancy. Loss of an MP Router or loss of connectivity through one of the IP networks does not disrupt the fabric connections.

**Figure 119 Redundant MP Routers and FCIP connections**



## C-series MDS 9216i, 9222i, IPS-4, IPS-8, 14/2 Multiprotocol Services Modules, 18/4 Multiservice Modules

The C-series MDS 9216i, MDS 9222i, IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Services Modules, and 18/4 Multiservice Modules provide MDS FCIP and iSCSI functionality.

The IPS-4, IPS-8, 14/2, and 18/4 modules integrate seamlessly into the C-series MDS 9000 switches and support the full range of features, including VSANs, security, and traffic management.

You can use the C-series modules in the C-series 9500 and 9200 series switches. The IPS-4, IPS-8, 14/2, and 18/4 modules have four, eight, two, and four 1 Gb/s Ethernet ports, respectively.

Table 160 (page 287) lists the features and requirements for the C-series modules.

**Table 160 C-series MDS module features and requirements**

Feature	Requirements
Fibre Channel switch hardware support for iSCSI and FCIP with the IP services modules	<ul style="list-style-type: none"> <li>C-series MDS 9216/9216A/9216i Multilayer Fabric switches</li> <li>C-series MDS 9222i Multiservice Fabric</li> <li>C-series MDS 9513 Multilayer Director switch</li> <li>C-series MDS 9506 Multilayer Director switch</li> </ul>

**Table 160 C-series MDS module features and requirements** (continued)

Feature	Requirements
	C-series MDS 9509 Multilayer Director switch
IP network protocols	TCP/IP IPv6, Ethernet 10 Mb/s, 100 Mb/s, and 1,000 Mb/s Requires dedicated IP bandwidth, see <a href="#">Table 164 (page 291)</a> and <a href="#">Table 168 (page 294)</a> .
Storage systems	HP P6000 Continuous Access HP P9000 (XP) Continuous Access Contact an HP storage representative for EVA and XP supported models. For specific operating system support, see “Heterogeneous server rules” (page 159). Supported load balance settings: <ul style="list-style-type: none"> <li>HP P9000 (XP) Continuous Access—Supported with load balance setting <code>src-dst-id</code> or <code>src-dst-ox-id</code> on all C-series Fibre Channel switches.</li> <li>HP P6000 Continuous Access—Supported with load balance setting <code>src-dst-id</code> or <code>src-dst-ox-id</code> (XCS 09534000 or later) on all C-series Fibre Channel switches.</li> <li>For the MDS 9020, the default setting is <code>src-dst-id</code>.</li> </ul>
FCIP-supported operating systems	For HP P6000 Continuous Access, see the <i>HP P6000 Enterprise Virtual Array Compatibility Reference</i> at <a href="http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html">http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html</a> .
Documentation	For information about the C-series MDS modules, see <a href="http://www.cisco.com/en/US/products/hw/ps4159/ps4358/products_data_sheet09186a00800c465b.html">http://www.cisco.com/en/US/products/hw/ps4159/ps4358/products_data_sheet09186a00800c465b.html</a> .

## HP storage replication products

HP provides the following storage replication products:

- HP P6000 Continuous Access with the P6300/P6500 EVA or EVA4000/4100/4400/6000/6100/6400/8000/8100/8400 using XCS
- HP P6000 Continuous Access with the EVA3000/5000 using VCS
- HP P9000 (XP) Continuous Access with the XP family of storage systems
- OpenVMS host-based volume shadowing

The following products are qualified by HP as Fibre Channel routers, network gateways, or iSCSI bridges with the EVA storage system:

- 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, MP Router Blade, and MP Router are qualified as FCIP gateways and Fibre Channel routers. The 1606 Extension SAN Switch, DC Dir Switch MP Extension Blade, 400 MP Router, and MP Router Blade are also qualified as Fibre Channel switches.
- C-series IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Services Module, MDS 9216i, and MDS 9222i switches are qualified as an FCIP gateway and an iSCSI bridge. MDS 9216i and MDS 9222i switches are also qualified as Fibre Channel switches.
- MPX200 Multifunction Router is qualified as an FCIP gateway.
- HP IP Distance Gateway is qualified as an FCIP gateway.

The following sections describe network requirements for replication products with qualified gateways:

- “SAN extension best practices for HP P6000 Continuous Access” (page 289)
- “HP P6000 Continuous Access with XCS 10x or XCS 09x” (page 289)
- “HP P6000 Continuous Access with XCS 6.x” (page 291)
- “HP P6000 Continuous Access with VCS 4.x” (page 293)



- “HP P9000 (XP) Continuous Access” (page 296)
- “OpenVMS host-based volume shadowing” (page 310)

## SAN extension best practices for HP P6000 Continuous Access

HP recommends you consider the following best practices when implementing SAN extension using FCIP with HP P6000 Continuous Access:

- Separate host traffic from replication traffic in environments that employ FCIP gateways. This decouples the host I/O from throughput on the inter-site link. This can be achieved in the following ways:
  - Use separate switches and separate fabrics for host and replication I/O (preferred solution).
  - Replication zones can be set up through fabric zoning. Using standard fabric zoning, the host traffic is separated from the replication traffic.
- For high availability and no single point of failure, the best practice is to deploy a 6-fabric solution, where there are two inter-site links. This solution uses separate switches for the replication and host fabrics. At least one of the EVAs must be an EVA8x00. Two fabrics that are dedicated to replication and four fabrics are dedicated to host I/O traffic.
- If only a single ISL is available, then a 5-fabric configuration can be used. This solution requires using separate switches for the host and replication fabrics. Only one fabric would be dedicated to replication traffic.
- Use a 2-fabric (dual fabric) configuration only if the host traffic cannot be separated from the replication traffic. This is recommended only when the customer business requirements permit it. Take care when designing a 2-fabric configuration over FCIP links, as host port blocking can occur.

For information on the different fabric configurations or on the HP SCSI Fibre Channel protocol option, see the *HP P6000 Continuous Access Implementation Guide*.

## HP P6000 Continuous Access with XCS 10x or XCS 09x

This section describes the HP P6000 Continuous Access with P6300/P6500 EVA XCS 10x and EVA4400/6400/8400 10x or XCS 09x data replication specifications and the supported minimum and maximum transmission rates for qualified switch and gateway pairs. [Table 161 \(page 289\)](#) lists the network requirements for long-distance gateways when using XCS 10x and XCS 09x.

**Table 161 Network requirements for long-distance IP gateways with XCS 10x or XCS 09x**

Specification	Description
IP bandwidth <sup>1</sup>	<ul style="list-style-type: none"> <li>• Must be dedicated to the HP P6000 Continuous Access storage replication function.</li> <li>• The minimum IP bandwidth required for HP P6000 Continuous Access with P6300/P6500 EVA, EVA4400/6400/8400, and EVA4400 (embedded switch) with FCIP is 2 Mb/s per path, or 4 Mb/s for two paths when using one IP link.</li> <li>• There is no support for dynamic pacing of the gateway.</li> </ul>
Maximum number of DR groups	<p>For the maximum number of DR groups, see <a href="#">Table 130 (page 220)</a>.</p> <p>For minimum supported bandwidth and resulting maximum number of DR groups based on the average packet-loss ratio and one-way inter-site latencies, see <a href="#">Table 162 (page 290)</a> and <a href="#">Table 163 (page 291)</a>.</p>
MTU of the IP network	1,500 bytes
Maximum latency <sup>1</sup>	100 ms IP network delay one-way or 200 ms round-trip
Average packet-loss ratio <sup>2</sup>	Low-loss network: 0.0012% average over 24 hours

**Table 161 Network requirements for long-distance IP gateways with XCS 10x or XCS 09x (continued)**

Specification	Description
	High-loss network: 0.2% average over 24 hours; must not exceed 0.5% for more than 5 minutes in a two-hour window
Latency jitter <sup>3</sup>	Must not exceed 10 ms over 24 hours

<sup>1</sup> Pre-existing restriction

<sup>2</sup> A high packet-loss ratio indicates the need to retransmit data across the inter-site link. Each retransmission delays transmissions queued behind the current packet, thus increasing the time to complete pending transactions. Unless noted otherwise, gateways listed in Table 162 (page 290) and Table 163 (page 291) are supported in both low-loss and high-loss networks.

<sup>3</sup> Latency jitter is the difference between the minimum and maximum values, and indicates how stable or predictable the network delay. The greater the jitter, the greater the variance in the delay, which lowers the performance predictability.

**NOTE:** Applications typically require more than the minimum bandwidth to meet throughput requirements. To increase the maximum number of DR groups, you must increase the minimum available IP bandwidth. For example, if the maximum number of DR groups required is 10, increase the minimum available bandwidth to 10 Mb/s; for 15 DR groups, increase it to 15 Mb/s; and for 128 DR groups, increase it to 128 Mb/s.

Table 162 (page 290) describes the network bandwidth requirements and IP gateway support for XCS 09x with B-series switches.

**Table 162 Network requirements for long-distance IP gateways and B-series switches with XCS 10x or XCS 09x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>	At least 2 Mb/s for 1 DR group	At least 4 Mb/s for 1 DR group
MPX200 Multifunction Router			
B-series 1606 Extension SAN Switch DC Dir Switch MP Extension Blade	See Table 22 (page 86).	Recommended: At least 5 Mb/s for 1 to 5 DR groups	Recommended: At least 10 Mb/s for 1 to 5 DR groups
B-series 400 MP Router	See Table 23 (page 87).		

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 163 (page 291) describes the network bandwidth requirements and IP gateway support for XCS 9.x with C-series switches.

**Table 163 Network requirements for long-distance IP gateways and C-series switches with XCS 09x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
MPX200 Multifunction Router			
C-series IPS-8, 18/4, MDS 9216i, MDS 9222i	See Table 38 (page 118).		

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

## HP P6000 Continuous Access with XCS 6.x

This section describes the HP P6000 Continuous Access XCS 6.x data replication specifications and the supported minimum and maximum transmission rates for qualified switch and gateway pairs. Table 164 (page 291) lists the network requirements for long-distance gateways that work with EVA XCS 6.x.

**Table 164 Network requirements for long-distance IP gateways with XCS 6.x**

Specification	Description
IP bandwidth <sup>1</sup>	<ul style="list-style-type: none"> <li>Must be dedicated to the HP P6000 Continuous Access storage replication function.</li> <li>The minimum IP bandwidth required for HP P6000 Continuous Access with FCIP is 2 Mb/s per path, or 4 Mb/s for two paths when using one IP link.</li> <li>There is no support for dynamic pacing of the gateway.</li> </ul>
Maximum number of DR groups	For the maximum number of DR groups, see Table 130 (page 220). for minimum supported bandwidth and resulting maximum number of DR groups based on the average packet-loss ratio and one-way inter-site latencies, see Table 165 (page 292) and Table 166 (page 293).
MTU of the IP network	1,500 bytes
Maximum latency <sup>1</sup>	100 ms IP network delay one-way or 200 ms round-trip
Average packet-loss ratio <sup>2</sup>	Low-loss network: 0.0012% average over 24 hours High-loss network: 0.2% average over 24 hours; must not exceed 0.5% for more than 5 minutes in a two-hour window
Latency jitter <sup>3</sup>	Must not exceed 10 ms over 24 hours

<sup>1</sup> Pre-existing restriction

<sup>2</sup> A high packet-loss ratio indicates the need to retransmit data across the inter-site link. Each retransmission delays transmissions queued behind the current packet, thus increasing the time to complete pending transactions. Unless noted

otherwise, gateways listed in [Table 165 \(page 292\)](#) and [Table 166 \(page 293\)](#) are supported in both low-loss and high-loss networks.

<sup>3</sup> Latency jitter is the difference between the minimum and maximum values, and indicates how stable or predictable the network delay. The greater the jitter, the greater the variance in the delay, which lowers the performance predictability.

**NOTE:** Applications typically require more than the minimum bandwidth to meet throughput requirements. To increase the maximum number of DR groups, you must increase the minimum available IP bandwidth. For example, if the maximum number of DR groups required is 10, increase the minimum available bandwidth to 10 Mb/s; for 15 DR groups, increase to 15 Mb/s; and for 128 DR groups, increase the minimum available bandwidth to 128 Mb/s.

[Table 165 \(page 292\)](#) describes the network bandwidth requirements and IP gateway support for XCS 6.x with B-series switches.

**Table 165 Network requirements for long-distance IP gateways and B-series switches with XCS 6.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>		
MPX200 Multifunction Router			
B-series 1606 Extension SAN Switch DC Dir Switch MP Extension Blade	See <a href="#">Table 22 (page 86)</a> .	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
B-series 400 MP Router, MP Router Blade	See <a href="#">Table 23 (page 87)</a> .		
<b>Legacy IP gateways</b>			
B-series MP Router	See <a href="#">Table 23 (page 87)</a> .	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
HP SR2122-2	3.4.5c	At least 10 Mb/s for 1 to 10 DR groups	At least 20 Mb/s for 1 to 10 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see [Table 130 \(page 220\)](#).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

[Table 166 \(page 293\)](#) describes the network bandwidth requirements and IP gateway support for XCS 6.x with C-series switches.

**Table 166 Network requirements for long-distance IP gateways and C-series switches with XCS 6.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
MPX200 Multifunction Router			
C-series IPS-8, 14/2, 18/4, MDS 9216i, MDS 9222i	See Table 38 (page 118).		
<b>Legacy IP gateways</b>			
C-series IPS-4	See Table 38 (page 118).	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
HP SR2122-2	3.4.5c	At least 10 Mb/s for 1 to 10 DR groups	At least 20 Mb/s for 1 to 10 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 167 (page 293) describes the network bandwidth requirements and IP gateway support for XCS 6.00 with M-series switches.

**Table 167 Network requirements for long-distance IP gateways and M-series switches with XCS 6.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
McDATA Eclipse 1620 SAN Router, Eclipse 2640 SAN Router	EOSi 4.6	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

## HP P6000 Continuous Access with VCS 4.x

This section describes the VCS 4.x data replication specifications and the supported minimum and maximum transmission rates for qualified switch and gateway pairs.

Table 168 (page 294) lists the network requirements for long-distance IP gateways that work with EVA VCS 4.x.

**Table 168 Network requirements for long-distance IP gateways with VCS 4.x**

Specification	Description
IP bandwidth <sup>1</sup>	<ul style="list-style-type: none"><li>• Must be dedicated to the HP P6000 Continuous Access storage replication function.</li><li>• The minimum IP bandwidth required for HP P6000 Continuous Access with FCIP is 2 Mb/s per path, or 4 Mb/s for two paths when using one IP link.</li><li>• There is no support for dynamic pacing of the gateway.</li></ul>
Maximum number of DR groups	For the maximum number of DR groups, see <a href="#">Table 130 (page 220)</a> . For minimum supported bandwidth and resulting maximum number of DR groups based on the average packet-loss ratio and one-way inter-site latencies, see <a href="#">Table 169 (page 295)</a> , <a href="#">Table 170 (page 296)</a> , and <a href="#">Table 171 (page 296)</a> .
MTU of the IP network	1,500 bytes
Maximum latency <sup>1</sup>	100 ms one-way or 200 ms round-trip
Average packet-loss ratio <sup>2</sup>	Low-loss network: 0.0012% average over 24 hours High-loss network: 0.2% average over 24 hours; must not exceed 0.5% for more than 5 minutes in a two-hour window
Latency jitter <sup>3</sup>	Must not exceed 10 ms over 24 hours

<sup>1</sup> Pre-existing restriction

<sup>2</sup> A high packet-loss ratio indicates the need to retransmit data across the inter-site link. Each retransmission delays transmissions queued behind the current packet, thus increasing the time to complete pending transactions. Unless noted otherwise, gateways listed in [Table 169 \(page 295\)](#) and [Table 170 \(page 296\)](#) are supported in both low-loss and high-loss networks.

<sup>3</sup> Latency jitter is the difference between the minimum and maximum values, and indicates how stable or predictable the network delay. The greater the jitter, the greater the variance in the delay, which lowers the performance predictability.

**NOTE:** Applications typically require more than the minimum bandwidth to meet throughput requirements. To increase the maximum number of DR groups, you must increase the minimum available IP bandwidth. For example, if the maximum number of DR groups required is 10, increase the minimum available bandwidth to 10 Mb/s; for 15 DR groups, increase to 15 Mb/s; and for 128 DR groups, increase the minimum available bandwidth to 128 Mb/s.

Table 169 (page 295) describes the network bandwidth requirements and IP gateway support for VCS 4.x with B-series switches.

**Table 169 Network requirements for long-distance IP gateways and B-series switches with VCS 4.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>		
MPX200 Multifunction Router			
B-series 1606 Extension SAN Switch DC Dir Switch MP Extension Blade	See Table 22 (page 86).	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups
B-series 400 MP Router, MP Router Blade	See Table 23 (page 87).		
<b>Legacy IP gateways</b>			
B-series MP Router	See Table 23 (page 87).	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 170 (page 296) describes the network bandwidth requirements and IP gateway support for VCS 4.x with C-series switches.

**Table 170 Network requirements for long-distance IP gateways and C-series switches with VCS 4.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
HP IP Distance Gateway (mpx110)	See note. <sup>3</sup>	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups
MPX200 Multifunction Router			
C-series IPS-8, 14/2, 18/4, MDS 9216i, MDS 9222i	See Table 38 (page 118).		
<b>Legacy IP gateways</b>			
C-series IPS-4	See Table 38 (page 118).	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 4 Mb/s for 1 DR group Recommended: At least 10 Mb/s for 1 to 5 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

<sup>3</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 171 (page 296) describes the network bandwidth requirements and IP gateway support for VCS 4.x with M-series switches.

**Table 171 Network requirements for long-distance IP gateways and M-series switches with VCS 4.x**

Gateway pair	Minimum supported firmware version	Minimum IP bandwidth <sup>1</sup> and maximum DR groups <sup>2</sup>	
		Dual IP link maximum latency	Single or shared IP link maximum latency
		0 to 100 ms one-way	0 to 100 ms one-way
McDATA Eclipse 1620 SAN Router, Eclipse 2640 SAN Router	EOSi 4.6	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups	At least 2 Mb/s for 1 DR group Recommended: At least 5 Mb/s for 1 to 5 DR groups

<sup>1</sup> HP P6000 Continuous Access requires a minimum of 2 Mb/s of IP bandwidth per path, or 4 Mb/s for two paths.

<sup>2</sup> Assumes single-member DR groups (1 virtual disk). For the maximum number of DR groups supported based on maximum IP bandwidth, see Table 130 (page 220).

## HP P9000 (XP) Continuous Access

HP P9000 (XP) Continuous Access is supported on all XP storage systems. It supports three data replication modes: Synchronous, Asynchronous, and Journal. Table 172 (page 297) describes the XP storage systems, replication modes, and required firmware versions.

HP P9000 (XP) Continuous Access Synchronous supports two fence levels: data and never (status fence level is not supported).



- **Data**—Prevents host writes in the event of a replication link failure
- **Never**—Allows host writes to continue after a replication link failure

**Table 172 HP P9000 (XP) Continuous Access replication modes<sup>1</sup>**

Storage system	Synchronous	Asynchronous <sup>2</sup>	Journal <sup>3</sup>	Recommended firmware versions	Minimum firmware versions
P9500 <sup>4</sup>	•	—	•	V01	V01+1
XP24000 XP20000	•	•	•	60.06.05.00/00	60.01.68.00/00
XP12000 XP10000	•	•	•	50.09.86.00/00 <sup>5</sup>	50.05.46.00/00 <sup>5</sup>
XP1024 XP128	•	•	—	21.14.50.00/00 <sup>6</sup>	21.13.02.00/00 <sup>6</sup>
XP512 XP48	•	•	—	01.19.99.00/02	01.19.86.00/00

<sup>1</sup> **Legend:** • = supported; — = not supported

<sup>2</sup> XP24000/20000/12000/10000: 32K pairs

<sup>3</sup> XP24000 with RAID Manager 1.20.05 (or later), XP12000/10000 with RAID Manager 1.17.04 (or later), P9500 with RAID Manager 1.24.16 (or later)

<sup>4</sup> 64K pairs

<sup>5</sup> Replication between XP24000/20000 and XP12000/10000 requires firmware 50.09.37.00/02 (or later).

<sup>6</sup> Replication between XP24000/20000 and XP1024/128 requires firmware 21.14.37.00/00 (or later).

HP P9000 (XP) Continuous Access is supported using the Fibre Channel protocol or ESCON protocol. Current P9500 and XP storage systems support replication using Fibre Channel. Legacy XP storage systems use either Fibre Channel or ESCON. Data replication between different storage systems is supported only if they use the same protocol.

Table 173 (page 297) describes the supported protocol and storage system source-target combinations.

**Table 173 HP P9000 (XP) Continuous Access protocols and source-target replication pairing<sup>1</sup>**

Storage system	Protocol	P9500	XP24000 XP20000	XP12000 XP10000	XP1024 XP128		XP512 XP48	
		Fibre Channel	Fibre Channel	Fibre Channel	Fibre Channel	ESCON	Fibre Channel	ESCON
P9500	Fibre Channel	•	•	•	—	—	—	—
XP24000 XP20000	Fibre Channel	•	•	•	•	—	—	—
XP12000 XP10000	Fibre Channel	•	•	•	•	—	•	—
XP1024 XP128	Fibre Channel	—	•	•	•	—	•	—
	ESCON	—	—	—	—	•	—	•

**Table 173 HP P9000 (XP) Continuous Access protocols and source-target replication pairing** *(continued)*

Storage system	Protocol	P9500	XP24000 XP20000	XP12000 XP10000	XP1024 XP128		XP512 XP48	
		Fibre Channel	Fibre Channel	Fibre Channel	Fibre Channel	ESCON	Fibre Channel	ESCON
XP512 XP48	Fibre Channel	—	—	•	•	—	•	—
	ESCON	—	—	—	—	•	—	•

<sup>1</sup> Legend: • = supported; — = not supported

When performing replication between two storage systems, the following restrictions apply:

- Legacy storage systems have limitations on the control-unit range, port numbers, and LUNs.
- Support for “host group” is limited by storage system functionality. This also affects the internal addressing of devices allocated to host ports and the configuration of RAID Manager.
- The emulation mode and the device (volume) size must be the same on both storage systems.
- There are firmware requirements for replication between different P9500 and XP storage systems. For firmware versions, see [Table 172 \(page 297\)](#).

### HP P9000 (XP) Continuous Access Synchronous

[Table 174 \(page 298\)](#) describes the network requirements and configuration rules when using HP P9000 (XP) Continuous Access Synchronous replication.

**Table 174 HP P9000 (XP) Continuous Access Synchronous replication rules**

Rule number	Description
1	Synchronous replication is supported on P9500, XP24000/20000, XP12000/10000, XP1024/128, and XP512/48 storage systems. The maximum distance or latency supported for synchronous replication is 200 km when using dark fiber, or 20 ms using other extension methods. Latencies greater than 20 ms require HP approval before implementation. <b>Note:</b> You must ensure that host-based applications utilizing replicated storage are able to tolerate the total of the following latencies: <ul style="list-style-type: none"> <li>• Local I/O servicing (local site storage system)</li> <li>• Continuous Access Synchronous replication</li> <li>• Remote I/O servicing (remote site storage system)</li> </ul>
2	A minimum of 16 Mb/s of IP bandwidth per path is required.

### HP P9000 (XP) Continuous Access Asynchronous

[Table 175 \(page 298\)](#) describes the network requirements and configuration rules when using HP P9000 (XP) Continuous Access Asynchronous replication.

**Table 175 HP P9000 (XP) Continuous Access Asynchronous replication rules**

Rule number	Description
1	Asynchronous replication is supported on XP24000/20000, XP12000/10000, XP1024/128, and XP512/48 storage systems. The maximum latency supported for Asynchronous replication is 300 ms one-way or 600 ms round-trip. <b>Note:</b> P9500 storage systems do not support Continuous Access Asynchronous replication.
2	A minimum of 16 Mb/s of IP bandwidth per path is required.

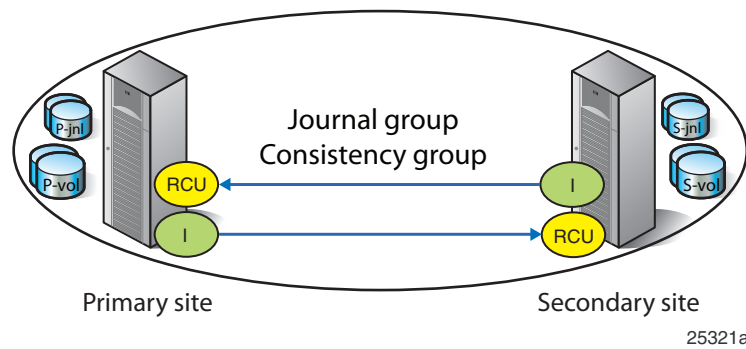
## HP P9000 (XP) Continuous Access Journal

Table 176 (page 299) describes the network requirements and configuration rules when using HP P9000 (XP) Continuous Access Journal replication.


**Table 176 HP P9000 (XP) Continuous Access Journal replication rules**

Rule number	Description
1	Journal replication is supported on and P9500 with V01, XP24000/20000 and XP12000/XP10000 storage systems only. The maximum distance supported is 300 ms one-way or 600 ms round-trip.
2	Initiator port and RCU-target port configurations and a control-unit-free definition are required on both storage systems. See <a href="#">Figure 120 (page 299)</a> .
3	Open-V emulation is supported for data volumes (P-vol and S-vol) and journal pools only.
4	Journal replication does not support LUSE devices for P-vol, S-vol, or journal pools for firmware versions earlier than 50.09.07.00/00.
5	P-vol and S-vol must have the same size Open-V volumes.
6	P-jnls and S-jnls must have at least 1 LDEV and can have up to 16 Open-V LDEVs. See <a href="#">Figure 121 (page 300)</a> .
7	HP recommends that the P-jnls and S-jnls in a journal group be the same size and have the same number and type of LDEVs.
8	An XP 3DC configuration requires minimum firmware versions for XP24000/20000 and XP12000/10000: <ul style="list-style-type: none"> <li>Two topologies are supported; cascaded (1:1:1) and multi-target (1:2). See <a href="#">Figure 122 (page 300)</a>.</li> </ul> XP12000/10000 firmware earlier than 50.08.05.00/00 did not support creating Business Copy from the devices used for both Synchronous and Journal replication in a 3DC configuration. Firmware version 50.08.05.00/00 (or later) allows creation of Business Copy from these devices, but does not allow a fast-restore operation for a Business Copy device.
9	A P-vol can support a maximum of two remote copies (one Continuous Access Journal copy and one Continuous Access Synchronous copy).
10	Requires a minimum guaranteed network bandwidth link that matches the average host writes of all MCU journals being serviced. <p>Journal capacity must be capable of buffering data until off-peak times.</p> <ul style="list-style-type: none"> <li>For 256 Mb/s or higher, set the line speed to 256.</li> <li>For 100 to 256 Mb/s, set the line speed to 100.</li> <li>For 10 to 100 Mb/s, set the line speed to 10.</li> </ul> <p><b>Note:</b> A minimum of 16Mb/s of IP bandwidth per path is required.</p>

**Figure 120 HP P9000 (XP) Continuous Access Journal configuration**

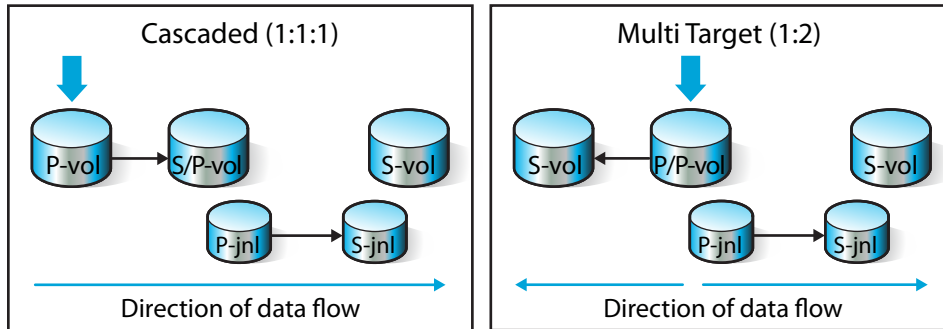


**Figure 121 HP P9000 (XP) Continuous Access P-jnl and S-jnl groups**



25322a

**Figure 122 HP P9000 (XP) Continuous Access 3DC configuration**



25339a

### HP P9000 (XP) Continuous Access configuration support


This section describes the products and maximum distances supported for the following HP P9000 (XP) Continuous Access configuration types:

- Direct storage-to-storage
- Fibre Channel switches
- ESCON directors and repeaters
- WDM
- ATM and SONET/SDH
- FCIP and routing extension

### Direct storage-to-storage

Figure 123 (page 301) shows an HP P9000 (XP) Continuous Access direct storage-to-storage configurations for Fibre Channel and ESCON. Table 177 (page 301) describes the maximum supported distances.

**Figure 123 HP P9000 (XP) Continuous Access direct storage-to-storage configurations for Fibre Channel and ESCON**




**Table 177 HP P9000 (XP) Continuous Access direct storage-to-storage distances**

Configuration	Maximum distance
Fibre Channel (Fibre CHIP) Port parameters for direct connect: <ul style="list-style-type: none"> <li>Fabric = Off, Connection = FC-AL</li> </ul>	150 m at 4 Gb/s 300 m at 2 Gb/s 500 m at 1 Gb/s
	Long-wave SFPs: 10 km at 4 Gb/s
ESCON	3 km

**Fibre Channel switches**

Figure 124 (page 301) shows an HP P9000 (XP) Continuous Access Fibre Channel configuration. Table 178 (page 302) describes the maximum supported distances.

**Figure 124 HP P9000 (XP) Continuous Access single-switch and multi-switch Fibre Channel configurations**




**Table 178 HP P9000 (XP) Continuous Access Fibre Channel distances**

Configuration	Maximum distance
<p>Single-switch or multi-switch See Figure 124 (page 301).</p> <p><b>Note:</b> For Fibre Channel switch model support, contact an HP storage representative.</p> <p>For B-series switches and routers:</p> <ul style="list-style-type: none"> <li>• Supports port-based routing (aptpolicy = 1) and exchange-based routing (aptpolicy = 3). See also Table 183 (page 307).</li> </ul> <p>Port parameters for switch/fabric connect:</p> <ul style="list-style-type: none"> <li>• Fabric = On, Connection = Point to Point (recommended setting)</li> <li>• Fabric = On, Connection = FC-AL</li> </ul> <p><b>Note:</b> For all switch configurations, HP P9000 (XP) Continuous Access ports and host ports must be in separate zones.</p>	<p>See Table 76 (page 150), Table 77 (page 150), and Table 80 (page 151).</p> <p><b>Note:</b> For WDM distances, see Table 180 (page 303).</p>


ESCON directors and repeaters

Figure 125 (page 302) and Figure 126 (page 302) show HP P9000 (XP) Continuous Access ESCON director and repeater configurations. Table 179 (page 302) describes the maximum supported distances.

**Figure 125 HP P9000 (XP) Continuous Access ESCON director configuration**



**Figure 126 HP P9000 (XP) Continuous Access ESCON repeater configuration**



**Table 179 HP P9000 (XP) Continuous Access ESCON director and repeater distances**

Configuration	Maximum distance
IBM 9032/9033 director	3 km for short-wave, multi-mode ESCON (MME)
IBM 9036 repeater	20 km for long-wave, single-mode ESCON (SME)
Nbase Xyplex	43 km for director/repeater combination

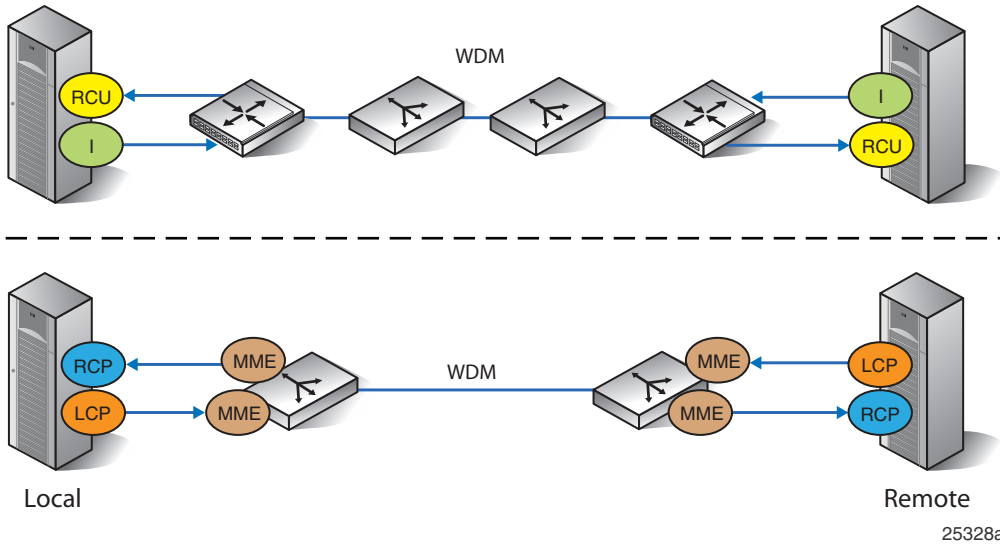
**Table 179 HP P9000 (XP) Continuous Access ESCON director and repeater distances** *(continued)*

Configuration	Maximum distance
<b>Note:</b> Supported only in a director-to-director configuration.	<b>Note:</b> The ESCON protocol is not converted in this process.

**WDM**

Figure 127 (page 303) shows HP P9000 (XP) Continuous Access WDM (CWDM and DWDM) configurations. Table 180 (page 303) describes the maximum supported distances and equipment.

**Figure 127 HP P9000 (XP) Continuous Access WDM configurations**




**Table 180 HP P9000 (XP) Continuous Access WDM distances and equipment**

Configuration	Maximum distance and equipment
Fibre Channel	For supported distances, see Table 76 (page 150), Table 77 (page 150), and Table 80 (page 151). <b>Note:</b> For Synchronous replication, distance impacts performance. HP recommends a maximum distance of 200 km when using dark fiber. HP P9000 (XP) Continuous Access is supported with all WDM products supported by the switch vendors. See "Certified third-party WDM products" (page 311).
ESCON	50 km Nortel Optera Metro 5200/5100 Movaz RAYexpress

**ATM and SONET/SDH**

Figure 128 (page 304) shows HP P9000 (XP) Continuous Access ATM and SONET/SDH configurations. Table 181 (page 304) describes the supported products.

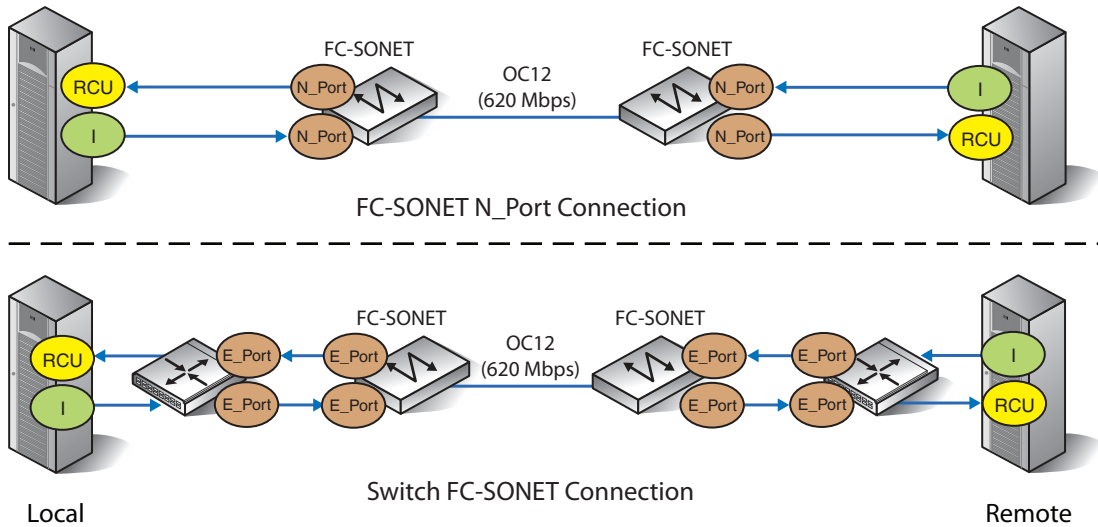
**Figure 128 HP P9000 (XP) Continuous Access ATM and SONET/SDH configurations**



25329a

Figure 129 (page 304) shows HP P9000 (XP) Continuous Access SONET/SDH FC direct and switch configurations.

**Figure 129 HP P9000 (XP) Continuous Access SONET/SDH FC direct and switch configurations**



25337b

**Table 181 HP P9000 (XP) Continuous Access ATM and SONET/SDH products<sup>1</sup>**

Configuration	XP storage system	Supported bandwidths and products
ESCON (ATM)	XP24000 XP20000	OC-3—Brocade-McDATA-CNT Ultraset Storage Director (USD) firmware 2.7 or 3.2.1
ESCON (ATM)	XP12000 XP10000	OC-3—Brocade-McDATA-CNT-Inrange 9801 SNS firmware 2.3 (Build 27) or 2.4 (Build 13) Brocade-McDATA-CNT-Inrange 9811H FW ACP-3; 3.1.1702 <b>Requirements:</b> <ul style="list-style-type: none"> <li>• Buffer size = 64 credits/port</li> <li>• 1 x ATM OC-3 (155 Mb/s) port</li> </ul>



**Table 181 HP P9000 (XP) Continuous Access ATM and SONET/SDH products** *(continued)*

Configuration	XP storage system	Supported bandwidths and products
		<ul style="list-style-type: none"> <li>Interoperability mode = 1</li> <li>E_D_TOV = 5,000</li> </ul> <p><b>Note:</b> This converter has been discontinued. There is limited support for existing installations.</p>
Fibre Channel	XP512 XP48	OC-3—Brocade-McDATA-CNT Edge 1003 with B-series switches OC-12—Brocade-McDATA-CNT Edge 1002 with B-series switches <b>Requirements:</b> <ul style="list-style-type: none"> <li>With Brocade-McDATA-CNT Edge firmware 1.5.1 and B-series switches, ISL R_RDY mode = ON must be set on the E_Port of the switch (portcfgISLMode = 1).</li> <li>A remote switch license is required for B-series switches using firmware 2.x.</li> <li>E_D_TOV = 5,000</li> <li>R_A_TOV = 10,000</li> <li>DataFieldSize = 2,048</li> </ul>
Fibre Channel	XP12000 XP10000 XP1024 XP128 XP512 XP48	OC-3—Brocade-McDATA-CNT (Inrange) 9801H or L, 2.3 (Build 28) or 2.4 (Build 13) <b>Requirements:</b> <ul style="list-style-type: none"> <li>Buffer size = 64 credits/port</li> <li>1 x ATM OC-3 (155 Mb/s) port</li> <li>Interoperability mode = 1</li> <li>E_D_TOV = 5,000</li> </ul> <p><b>Note:</b> This converter has been discontinued. There is limited support for existing installations.</p>
Fibre Channel	For a list of supported storage systems, contact your HP representative.	Cisco FCMR-4 (800-22030-03) for FC over SONET/SDH <b>Requirements:</b> <ul style="list-style-type: none"> <li>Firmware 57-5391-03, hardware revision 03</li> <li>FC/FICON (1 Gb/s or 2 Gb/s) over SONET/SDH (OC-3, OC-12, OC-48, or OC-192)</li> <li>255 buffer credits (ingress) and 1,200 buffer credits (egress)</li> <li>ML-Series, CE-Series, G-Series, or E-Series</li> </ul>
Fibre Channel	XP24000 XP20000 XP12000	Ciena CN 2000 for FC over SONET/SDH or DWDM. See <a href="#">Figure 129</a> (page 304). <b>Requirements:</b> <ul style="list-style-type: none"> <li>Firmware 5.1.0, 5.0.1</li> <li>SONET/SDH OC-3</li> <li>Port parameters:               <ul style="list-style-type: none"> <li>For direct connect, Fabric = Off, Connection = FC-AL</li> <li>For switch/fabric connect (B-series switches), Fabric = On, Connection = Point to Point</li> </ul> </li> </ul>
<b>Distances</b>		
<ul style="list-style-type: none"> <li>Distance is unlimited using converters and Asynchronous replication.</li> <li>Converter and Synchronous replication have practical distance limitations.</li> <li>Supported distance for an ATM connection depends on network latency delays, packet loss, and application performance requirements.</li> </ul>		


<sup>1</sup> Cisco FCMR-4 only

### FCIP and routing extension

This section describes the HP P9000 (XP) Continuous Access data replication IP specifications and the supported minimum and maximum transmission rates for qualified switch and IP gateway pairs.

Figure 130 (page 306) and Figure 131 (page 306) show HP P9000 (XP) Continuous Access FCIP and iFCP configurations.

**Figure 130 HP P9000 (XP) Continuous Access FCIP configuration**



**Figure 131 HP P9000 (XP) Continuous Access FCIP and iFCP configurations**

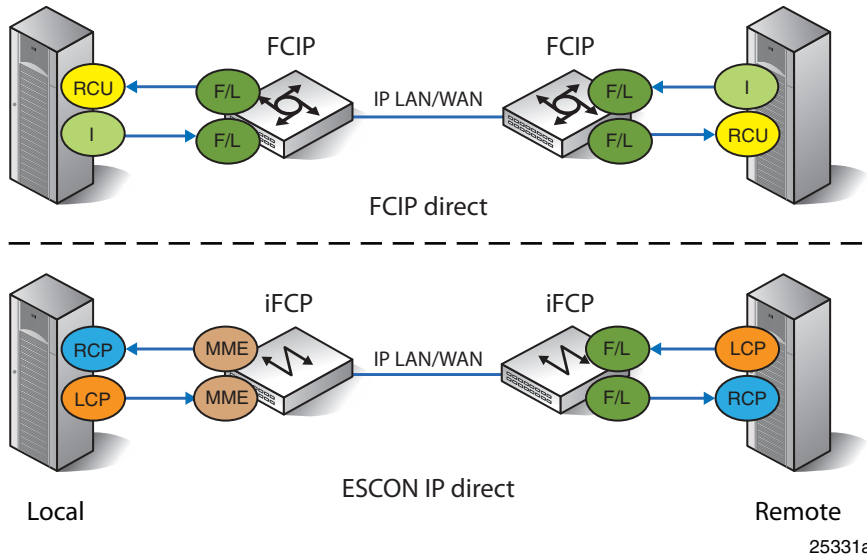


Table 182 (page 306) lists the network requirements for long-distance IP gateways that are supported with HP P9000 (XP) Continuous Access.

**Table 182 Network requirements for long-distance IP gateways with HP P9000 (XP) Continuous Access**

Specification	Requirement
Bandwidth <sup>1</sup>	<ul style="list-style-type: none"> <li>• Must be dedicated to the HP P9000 (XP) Continuous Access storage replication function.</li> <li>• The minimum IP bandwidth required for HP P9000 (XP) Continuous Access is 16 Mb/s per path.</li> <li>• When configuring multiple long-distance links, ensure all links provide equal bandwidth and latency for maximum sustained aggregate performance. If any single link has reduced performance or increased latency, all links will be reduced to the lowest performing link, significantly lowering the aggregate performance.</li> <li>• There is no support for dynamic pacing of the gateway.</li> </ul>
MTU of the IP network	<p>1,500 bytes: HP StorageWorks IP Distance Gateway (mpx100)                  2,348 bytes: B-series 400 MP Router, MP Router Blade                  3,000 bytes: C-series MDS 9222i, 9216i, IPS-8, IPS-4, 14/2</p> <p><b>Note:</b> These MTU settings are the recommended maximum values when using the respective FCIP products. You must ensure that the connected network components support the same values for end-to-end connectivity at the stated rates.</p>

**Table 182 Network requirements for long-distance IP gateways with HP P9000 (XP) Continuous Access** (continued)

Specification	Requirement
Maximum latency <sup>1</sup>	See Table 174 (page 298), Table 175 (page 298), and Table 176 (page 299).
Average packet-loss ratio <sup>2</sup>	Low-loss network: 0.0012% average over 24 hours High-loss network: 0.2% average over 24 hours; must not exceed 0.5% for more than 5 minutes in a two-hour window
Latency jitter <sup>3</sup>	Must not exceed 10 ms over 24 hours

<sup>1</sup> Pre-existing restriction

<sup>2</sup> A high packet-loss ratio indicates the need to retransmit data across the ISL. Each retransmission delays transmissions queued behind the current packet, thus increasing the time to complete pending transactions. Unless noted otherwise, gateways listed in Table 183 (page 307) are supported in both low-loss and high-loss networks.

<sup>3</sup> Latency jitter is the difference between the minimum and maximum values, and indicates how stable or predictable the network delay is. The greater the jitter, the greater the variance in the delay, which lowers the performance predictability.

**NOTE:** Applications typically require more than the minimum bandwidth to meet throughput requirements. For more information on link sizing, see the *HP StorageWorks XP Continuous Access User Guide*.

Table 183 (page 307), Table 184 (page 309), and Table 185 (page 310) list the IP gateways supported for use with HP P9000 (XP) Continuous Access.

**Table 183 HP P9000 (XP) Continuous Access FCIP gateway support**

Product and minimum supported firmware version	XP storage system	Notes
HP IP Distance Gateway (mpx110) <sup>1</sup>	P9500 XP24000 XP20000 XP12000 XP10000	Requires a B-series or C-series switch between the XP storage systems and the gateways with the following settings and packet-loss criteria: <ul style="list-style-type: none"> <li>• TCP window size = 5 or 6</li> <li>• Average packet-loss ratio maximum = 0.1%</li> </ul> Enable compression for IP fabrics with a RTT greater than or equal to 50 ms or guaranteed WAN bandwidth of less than or equal to 45 Mb/s. For performance tuning information based on the link speed and delay, see the <i>HP StorageWorks IP Distance Gateway User Guide</i> .
HP MPX200 Multifunction Router with FCIP <sup>1</sup>	P9500 XP24000 XP20000 XP12000 XP10000	Requires a B-series or C-series switch between the XP storage systems and the gateways with the following settings and packet-loss criteria: <ul style="list-style-type: none"> <li>• TCP window size = 5 or 6</li> <li>• Average packet-loss ratio maximum = 0.1%</li> </ul> For compression usage recommendations, see the <i>HP StorageWorks MPX200 Multifunction Router User Guide</i> at <a href="http://www.hp.com/go/mpx200">http://www.hp.com/go/mpx200</a> .
B-series 1606 Extension SAN Switch DC Dir Switch MP Extension Blade For firmware versions, see	P9500 XP24000 XP20000 XP12000	<ul style="list-style-type: none"> <li>• Can be used as a switch and FCIP gateway with XP storage systems directly connected</li> <li>• Can be used for Fibre Channel to Fibre Channel routing</li> <li>• Requires <code>iodSET</code> (in-order delivery)</li> <li>• Supports the following routing settings: Port-based routing:</li> </ul>

**Table 183 HP P9000 (XP) Continuous Access FCIP gateway support** *(continued)*

Product and minimum supported firmware version	XP storage system	Notes
Table 22 (page 86).	XP10000	<ul style="list-style-type: none"> <li>◦ aptpolicy = 1</li> <li>◦ aptpolicy -ap = 0 (port/asic load sharing)</li> <li>◦ dlsReset (no dynamic load sharing)</li> </ul> Exchange-based routing: <ul style="list-style-type: none"> <li>◦ aptpolicy = 3</li> <li>◦ aptpolicy -ap = 0 (port/asic load sharing)</li> <li>◦ dlsSet (dynamic load sharing)</li> </ul> For additional requirements, see Table 157 (page 281) and Table 178 (page 302).
B-series 400 MP Router and MP Router Blade For firmware versions, see Table 23 (page 87).	P9500 XP24000 XP20000 XP12000 XP10000 XP1024 XP128	<ul style="list-style-type: none"> <li>• Can be used as a switch and FCIP gateway with XP storage systems directly connected</li> <li>• Can be used for Fibre Channel to Fibre Channel routing</li> <li>• Maximum transmission rate from 16 Mb/s to 1 Gb/s</li> <li>• Requires ioDSET (in-order delivery)</li> <li>• Supports the following routing settings: Port-based routing:                             <ul style="list-style-type: none"> <li>◦ aptpolicy = 1</li> <li>◦ aptpolicy -ap = 0 (port/asic load sharing)</li> <li>◦ dlsReset (no dynamic load sharing)</li> </ul>                             Exchange-based routing:                             <ul style="list-style-type: none"> <li>◦ aptpolicy = 3</li> <li>◦ aptpolicy -ap = 0 (port/asic load sharing)</li> <li>◦ dlsSet (dynamic load sharing)</li> </ul> </li> </ul> For additional requirements, see Table 158 (page 283) and Table 178 (page 302).
B-series MP Router For firmware versions, see Table 22 (page 86).	XP24000 XP20000 XP12000 XP10000 XP1024 XP128	<ul style="list-style-type: none"> <li>• Requires a B-series switch between storage systems and the routers</li> <li>• Can be used for Fibre Channel to Fibre Channel routing</li> <li>• Maximum transmission rate from 16 Mb/s to 1 Gb/s</li> </ul> For additional requirements, see Table 159 (page 286).
C-series 9500 and 9200 series switches with IPS-8, IPS-4, 14/2, 18/4, MDS 9216i, and MDS 9222i For firmware versions, see Table 38 (page 118) and Table 41 (page 119).	P9500 XP24000 XP20000 XP12000 XP10000 XP1024 XP128 XP512 XP48	<ul style="list-style-type: none"> <li>• Can be used as a switch and FCIP gateway with XP storage systems directly connected</li> <li>• Distributed services TOV (D_S_TOV) = 5,000</li> <li>• Error detect TOV (E_D_TOV) = 2,000</li> <li>• Resource allocation TOV (R_A_TOV) = 10,000</li> <li>• DataFieldSize = 2,112</li> <li>• BB_credit from 16 (default) to 255, 14/2 and MDS9216i BB_credit maximum = 3,500</li> </ul>
HP SR2122-2 Firmware: 3.4.5c	XP12000 XP10000 XP1024 XP128	Supported switches: <ul style="list-style-type: none"> <li>• B-series</li> <li>• C-series</li> <li>• M-series</li> </ul>

**Table 183 HP P9000 (XP) Continuous Access FCIP gateway support (continued)**

Product and minimum supported firmware version	XP storage system	Notes
		For additional requirements, see Table 160 (page 287).

<sup>1</sup> For current support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

**Table 184 HP P9000 (XP) Continuous Access iFCP gateway support**

Product and minimum supported firmware version	XP storage system	Notes
Brocade-McDATA-CNT Edge 1000 and 1100 Firmware: 1.5.1.2.7	XP1024 XP128 XP512 XP48	<ul style="list-style-type: none"> <li>Firmware earlier than 1.5.1 requires a remote switch license.</li> <li>Supported switches: B-series 2400/2800, FC8B, and FC16B<sup>1</sup></li> <li>Data compression possible</li> </ul> <p><b>Note:</b> This converter has been discontinued. There is limited support for existing installations.</p>
Brocade-McDATA-CNT Edge 1001 and 1101 FC Firmware: 1.5.1.2.7	XP1024 XP128 XP512 XP48	<ul style="list-style-type: none"> <li>Firmware earlier than 1.5.1 requires a remote switch license.</li> <li>Supported switches: B-series 2400/2800, FC8B, and FC16B</li> <li>Data compression possible</li> </ul> <p><b>Note:</b> This converter has been discontinued. There is limited support for existing installations.</p>
Brocade-McDATA-CNT Edge 3000 Firmware: 3.1.1.3, 3.1.2, 3.1.4, 3.1.5	XP24000    XP1024 XP20000    XP128 XP12000    XP512 XP10000    XP48	<ul style="list-style-type: none"> <li>Supported switches: B-series 2/128, 2/64, 2/32, 2/16, 2800, and 2400<sup>1</sup></li> <li>Can be used in FCIP direct configuration</li> <li>Data compression possible</li> <li>No remote switch license required<sup>1</sup></li> <li>Maximum transmission rate from 16 Mb/s to 1 Gb/s</li> </ul>
McDATA-Nishan IPS3300	XP12000    XP128 XP10000    XP512 XP1024      XP48	Contact an HP storage representative.
McDATA-Nishan IPS4300	XP12000    XP128 XP10000    XP512 XP1024      XP48	Contact an HP storage representative.
McDATA Eclipse 2640 and 1620	XP24000    XP1024 XP20000    XP128	<ul style="list-style-type: none"> <li>Supported switches: McDATA 4400, 4700</li> <li>Can be used as a switch and iFCP gateway with XP storage systems</li> </ul>

**Table 184 HP P9000 (XP) Continuous Access iFCP gateway support** (continued)

Product and minimum supported firmware version	XP storage system	Notes
Firmware: 4.07 or 4.4.0	XP512 XP48 XP12000 XP10000 (1620 only for XP512/48)	directly connected
Brocade-McDATA-Inrange 9801 SNS Firmware: 2.3 (Build 28) or 2.4 (Build 13)	XP1024 XP128 XP512 XP48	<ul style="list-style-type: none"> <li>Supported switches: B-series 2400/2800, FC8B, and FC16B</li> <li>Requires a remote switch license</li> <li>Interoperability mode = 1, E_D_TOV = 5,000</li> </ul> <p><b>Note:</b> This converter has been discontinued. There is limited support for existing installations.</p>

<sup>1</sup> With McDATA-CNT Edge firmware 1.5.1 and B-series switches, ISL\_R\_RDY mode = ON must be set on the E\_Port of the switch (portcfgISLMode = 1). For B-series switches using firmware 2.x, a remote switch license is required. E\_D\_TOV = 5,000, R\_A\_TOV = 10,000, DataFieldSize = 2,048.

**Table 185 HP P9000 (XP) Continuous Access ESCON IP gateway support**

Configuration	Product and minimum supported firmware version	Minimum IP bandwidth per path
HP P9000 (XP) Continuous Access ESCON over IP	McDATA-CNT Ultraset Storage Director (USD) ESCON Firmware: 2.6.2-0 or 3.1	100 Mb/s

## OpenVMS host-based volume shadowing

HP supports the following HP and third-party devices and features for OpenVMS host-based volume shadowing:

- B-series 1606 Extension SAN Switch and DC Dir Switch MP Extension Blade
- 400 MP Router, MP Router Blade
- B-series FC and FCIP Fastwrite
- Cisco PA-FC-1G
- C-series MDS IP Storage Services Module (IPS-4, IPS-8)
- Cisco MDS 14/2 Multiprotocol Services Module (including in MDS 9216i Fabric Channel switch and gateway)
- Cisco MDS 18/4 Multiservice Module (including the MDS 9222i Multiservice Fabric switch and gateway)
- C-series Write Acceleration
- CNT UltraNet 1000
- CNT UltraNet 1001
- CNT UltraNet 3000
- SR2122-2 IP Storage Router

## Certified third-party WDM, iFCP, and SONET products

This section describes the following topics:

- “Certified third-party WDM products” (page 311)
- “Certified third-party iFCP and SONET products” (page 311)

### Certified third-party WDM products

HP supports HP P6000 Continuous Access, HP P9000 (XP) Continuous Access, and SVSP Synchronous Mirrors (RAID1) on all WDM products, including DWDM and CWDM, certified by the Fibre Channel switch vendors for the equivalent HP switch models.

---

**NOTE:** SVSP Continuous Access is only supported using native Fibre Channel connectivity.

---

- **B-series switch products**—All Brocade WDM-certified products, listed by Brocade switch or router model number, are supported. Contact an HP storage representative for the equivalent HP switch models. See the *Brocade Data Center Ready Compatibility Matrix* at <http://www.brocade.com/data-center-best-practices/resource-center/index.page> (select Matrices under All Available Resources).
- **C-series switch products**—All Cisco WDM-certified products, listed by Cisco switch model number, are supported. Cisco model numbers and HP model numbers are equivalent.
- **M-series switch products**—All McDATA WDM-certified products, listed by McDATA switch model number, are supported. Contact an HP storage representative for the equivalent HP switch models.

---

**NOTE:** HP CWDM support is based on the switch vendors' support for CWDM SFPs only. For information about HP-supported standard (non-CWDM) SFPs, see the HP switch model QuickSpecs or the Fibre Channel switch Streams on the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

For more information about WDM technology and product support, see “Wavelength division multiplexing” (page 253) and “WDM system architectures” (page 254).

---

### Certified third-party iFCP and SONET products

This section provides information on:

- iFCP products approved for use in HP P6000 Continuous Access environments (Table 186 (page 311))
- SONET products approved for use in HP P9000 (XP) Continuous Access environments (Table 181 (page 304))

**Table 186 Third-party iFCP products for SAN extension**

Model	Firmware version	Notes
McDATA Eclipse 1620 SAN Router	4.4	Supported only as an iFCP gateway
	EOSi 4.6	Certified on HP P6000 Continuous Access with M-series switches

For information about supported gateways for specific array firmware, see “HP storage replication products” (page 288).

---

## 18 iSCSI storage

This chapter describes iSCSI storage in an HP SAN environment:

- [“iSCSI overview”](#) (page 312)
- [“iSCSI concepts”](#) (page 312)
- [“iSCSI storage network requirements”](#) (page 316)
- [“HP Native iSCSI products”](#) (page 316)
- [“HP iSCSI bridge products”](#) (page 323)

### iSCSI overview

iSCSI is a storage transport protocol. The IETF developed iSCSI to encapsulate the SCSI protocol over an IP network. iSCSI has many of the same mechanisms as the Parallel SCSI and Fibre Channel protocols.

iSCSI facilitates creating SANs that include IP technology. iSCSI establishes and manages connections between IP-based hosts and storage systems. Many Fibre Channel switches and routers, as well as NAS systems, provide iSCSI support.

This section describes the following topics:

- [“iSCSI and Fibre Channel”](#) (page 312)
- [“iSCSI bridge to Fibre Channel”](#) (page 323)

iSCSI provides access to storage systems and SANs over standard Ethernet-based TCP/IP networks that can be dedicated to storage or in some cases, shared with traditional Ethernet applications.

---

**NOTE:** Existing TCP/IP networks might not support iSCSI storage. HP recommends using a dedicated Gigabit Ethernet network between iSCSI initiators and targets. This ensures adequate data security and performance. As an alternative, use IPsec to secure the connection on a public network with decreased performance. See the specific HP product requirements to determine if a dedicated IP network for storage is required.

---

### iSCSI and Fibre Channel

There are many factors to consider in choosing iSCSI or Fibre Channel, including:

- Lower deployment costs for iSCSI compared to Fibre Channel
- Widespread knowledge of IP technology for iSCSI
- Less expensive IP components
- iSCSI support for many server models
- Open architecture design with iSCSI
- iSCSI performance well matched to small and mid-range storage applications

### iSCSI concepts

This section describes key iSCSI concepts:

- [“Initiator and target devices”](#) (page 313)
- [“iSCSI naming”](#) (page 313)
- [“Discovery mechanisms”](#) (page 313)
- [“Sessions and logins”](#) (page 314)
- [“Security”](#) (page 315)



- “Software and hardware iSCSI initiators” (page 315)
- “Bridging and routing” (page 323)
- “iSCSI boot” (page 315)

## Initiator and target devices

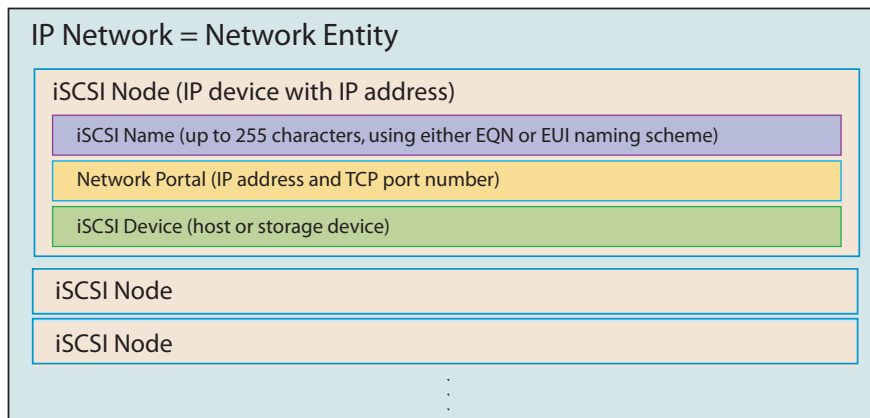
An iSCSI router manages access between iSCSI targets and iSCSI initiators as follows:

- **iSCSI target** (logical target)—An end-node device that is typically a storage system, storage router, or bridge. A storage system with iSCSI support is called native iSCSI storage.
- **iSCSI initiator** (IP host)—A system that starts the exchange of information with an iSCSI target. IP hosts access the iSCSI target storage systems as if they were directly attached.

## iSCSI naming

iSCSI nodes are uniquely named devices (initiators or targets). The nodes have an IP address, TCP port, and iSCSI name. The iSCSI name can be up to 255 characters in length. [Figure 132 \(page 313\)](#) shows an iSCSI node definition in the context of an IP network.

**Figure 132 iSCSI node definition**



25161a

The iSCSI name is independent of the network portal and provides a unique and consistent identity for an iSCSI node. Although moving a device to another network segment changes its network portal, the iSCSI name is unchanged and allows the device to be rediscovered.

An iSCSI name is independent of supporting hardware. You can assign an iSCSI name to a device driver on a host, even if the device driver accesses the network through multiple NICs. A storage device with multiple connections to the network is also identified by its iSCSI name.

iSCSI naming provides permanent and unique identities for iSCSI nodes. The two naming schemes are as follows:

- IQN
- EUI

An iSCSI node address consists of the IP address, TCP port number, and IQN or EUI. iSCSI nodes acquire IP addresses with standard IP services.

## Discovery mechanisms

This section describes the mechanisms you can use for discovery requests.

## Service Location Protocol

Clients (initiators) discover services (targets) using SLP, a client-server protocol. SLP for iSCSI uses three components:

- An iSCSI initiator has an SLP UA that serves as a client.
- iSCSI targets have an SLP SA that acts as an SLP server.
- A DA interprets multicast service requests from the server.

Initiators use three techniques for discovering targets:

- Unicast discovery service requests to the DA
- Multicast discovery service requests to SAs
- Unicast discovery service requests directly to an SA

## Static configuration

With static configuration, an administrator manually sets the target addresses for the initiators. The statically configured addresses for the targets persist across initiator reboots. HP recommends static configuration for the smallest iSCSI SANs.

## SendTargets command

With the `SendTargets` command, administrators configure the address of each target portal, setting up a range of target addresses for discovery. In a discovery session, an initiator sends the `SendTarget` command to discover all of the accessible target node names. HP recommends `SendTargets` for small iSCSI SANs.

## Internet Storage Name Service

The iSNS is a client-server discovery protocol. It provides naming and resource discovery services for storage systems on the IP network. The iSNS is modeled on both IP and Fibre Channel. iSNS components include:

- **iSNS server**—A directory server with optional security features.
- **iSCSI initiators with iSNS client capabilities**—The initiator iSNS client registers the initiator with the iSNS server and queries for a list of targets.
- **iSCSI targets with iSNS client capabilities**—The target iSNS client registers the target with the iSNS server.

## Sessions and logins

A session is a data exchange between an initiator and target. At the beginning of a session, information about the session is exchanged; later, application data is exchanged.

A session is enabled through an iSCSI login process:

1. The initiator establishes a TCP/IP connection.
2. The initiator starts the iSCSI login phase.
3. The initiator and target negotiate variable parameters.
4. Optional—The target verifies allowable connectivity with a security phase.
5. At the completion of the iSCSI login phase:
  - Success means the target sends a login accept to the initiator; the session continues.
  - Failure means the login is rejected; the TCP/IP connection is closed.

During iSCSI login, the initiator and target negotiate the lowest mutually acceptable value for each parameter. Negotiable parameters include:

- Type of security protocol, if any
- Maximum size of the data payload
- Support for unsolicited data
- Time-out values

During iSCSI login, the initiator and target also exchange nonnegotiable values such as names and aliases.

During an iSCSI session, unique session IDs are created for the initiator and target:

1. An initiator creates a unique ID by combining its iSCSI name with an ISID.
2. During login, the initiator sends the ISID to the target.
3. The target creates a unique ID by combining its iSCSI name with a TSID. The target sends the TSID to the initiator.

When login is complete, the iSCSI session enters the full-feature phase with normal iSCSI transactions.

## Security

Because iSCSI must accommodate untrusted IP environments, the specification for the iSCSI protocol defines multiple security methods:

- Encryption solutions that reside below the iSCSI protocol, such as IPsec, require no special negotiation between iSCSI end devices and are transparent to the upper layers.
- The iSCSI protocol has several encryption solutions including:
  - Kerberos
  - Public/private key exchanges

Security solutions can include an iSNS server that acts as a repository for public keys.

Text fields mediate the negotiation for the type of security supported by the end devices. If the negotiation is successful, the devices format their communications to follow the negotiated security routine.

## Software and hardware iSCSI initiators

An IP host can access an iSCSI environment using one of the following initiators:

- **Software iSCSI initiator**—The iSCSI code runs on the host and allows an Ethernet NIC to handle iSCSI traffic. Software iSCSI offers low cost with a performance penalty and CPU overhead. Software iSCSI initiators are available from many vendors.
- **TOE NIC**—Shifts processing of the communications protocol stack (TCP/IP) from the server processor to the NIC, lowering CPU overhead and use.
- **Hardware iSCSI initiator (iSCSI HBA)**—A high-performance HBA integrates both TCP/IP and iSCSI functions. Although integration adds cost to the HBA, it also provides high-speed iSCSI transport and minimal CPU overhead. The HBA transfers SCSI commands and data encapsulated by iSCSI directly to the host.

## iSCSI boot

iSCSI allows initiators (IP hosts) to boot from an iSCSI target. An iSCSI HBA typically has boot capabilities that must be enabled in its firmware.

## iSCSI storage network requirements

HP recommends:

- Dedicated IP network for iSCSI storage (can be required for some iSCSI products)
- Minimum GbE network bandwidth
- Multipathing driver when implementing high availability

See the iSCSI product sections for additional requirements:

- “HP StorageWorks Modular Smart Array 1510i” (page 317)
- “HP StorageWorks P4000 SAN” (page 321)
- “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324)
- “EVA and EVA4400 iSCSI Connectivity Option ” (page 333)
- “B-series MP Router and iSCSI” (page 344)
- “C-series iSCSI” (page 345)

## HP Native iSCSI products

This section describes HP native iSCSI storage system products. These products provide an iSCSI interface within the storage system hardware.

### P6300/P6500 EVA

- “P6300/P6500 EVA overview” (page 316)
- “P6300/6500 EVA iSCSI support” (page 316)

### P6300/P6500 EVA overview

The HP P6300/P6500 EVA storage systems are available with different SAN or host interface options:

- Fibre Channel interface only, four 8 Gb/s front end ports per controller
- Fibre Channel and iSCSI, two 8 Gb/s and four 1 GbE front end ports per controller
- Fibre Channel and iSCSI/FCoE, two 8 Gb/s and two 10-GbE front end ports per controller

This section describes iSCSI support. For information about P6300/P6500 EVA Fibre Channel support, see “P6000/EVA storage system rules” (page 213). For information about P6300/P6500 EVA FCoE support, see “Fibre Channel over Ethernet” (page 61).

### P6300/6500 EVA iSCSI support

The HP P6300/P6500 EVA iSCSI interface options provide 1 GbE iSCSI support (iSCSI module) or 10-GbE iSCSI support (iSCSI/FCoE module).

#### iSCSI or iSCSI/FCoE module

The iSCSI or iSCSI/FCoE modules are configured in a dual-controller configuration in the HP P6000. Dual-controller configurations provide for high availability with failover between iSCSI or iSCSI/FCoE modules. All configurations are supported as redundant pairs only. iSCSI connected servers can be configured for access to one or both controllers.

#### Operating system and multipath software support

For the latest information on operating system and multipath software support, see the HP SPOCKwebsite at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 187 (page 317) provides the operating system and multipath software support.

**Table 187 iSCSI/FCoE operating system and multipath software support**

Operating system	Multipath software	Clusters EVA storage system
Apple Mac OS X (1GbE iSCSI only)	None	None
Microsoft Windows Server 2008, 2003	MPIO with HP DSM MPIO with Microsoft DSM (with CN1000E only)	Failover Clustering, MSCS Hyper-V (1 GbE iSCSI only)
Red Hat Linux, SUSE Linux	Device Mapper	None
Solaris (1 GbE iSCSI only)	Solaris MPxIO	None
VMware	VMware MPxIO	None

For more information, see the iSCSI or iSCSI/FCoE configuration rules and guidelines chapter in the *HP P6300/P6500 Enterprise Virtual Array User Guide*.

## HP StorageWorks Modular Smart Array 1510i

This section describes the following topics:

- “MSA1510i overview” (page 317)
- “MSA1510i support” (page 317)

### MSA1510i overview

The HP StorageWorks Modular Storage Array 1510i (MSA1510i) is a controller shelf with an iSCSI interface that serves as an iSCSI target. Key features of the MSA1510i include:

- Supports MSA30 SCSI and MSA20 SATA (serial attachment) enclosures
- Two 1-GbE ports
- Path failover capability

For more information on the MSA1510i, see <http://h18006.www1.hp.com/products/storageworks/msa1510i>.

### MSA1510i support

The MSA1510i is compatible with many network devices and servers. It provides Windows operating system support and multiple management interfaces.

#### Network device support

The MSA1510i supports GbE components that conform to the IEEE 802.3ab (1000Base-T) standard. It is supported with Ethernet NICs that conform to this standard.

#### Server support

For the list of currently supported servers, see the product QuickSpecs at [http://h18000.www1.hp.com/products/quickspecs/12230\\_div/12230\\_div.html](http://h18000.www1.hp.com/products/quickspecs/12230_div/12230_div.html).

#### Operating system support

The MSA1510i supports the following Windows operating systems:

- Windows 2000 SP4 (32-bit)
- Windows Server 2003 (32-bit)

These operating systems require the Microsoft iSCSI initiator.

### Path failover software

The MSA1510i supports Microsoft MPIO basic failover software with an HP DSM driver. This software requires a second storage controller and iSCSI module.

### Management software support

The MSA1510i provides the web-based HP Storage Management Utility. It also provides a CLI, which uses a serial port, SSH, or Telnet. The MSA1510i supports SNMP.

## MSA1510i rules

Rules for using the MSA1510i include:

- Maximum of 20 simultaneous initiators
- Maximum of 32 LUNs

## HP StorageWorks MSA/2000 family of iSCSI SAN arrays

This section describes the following topics:

- “MSA2000i G2 and MSA2000i overview” (page 318)
- “P2000 G3 FC/iSCSI, MSA2000i G2 and MSA2000i support” (page 319)

### MSA2000i G2 and MSA2000i overview

The HP StorageWorks Modular Smart Array 2000i G2 (MSA2300i G2) and the HP StorageWorks Modular Smart Array 2000i (MSA2000i) are controller shelves with an iSCSI interface that serves as an iSCSI target. Key features include:

- Single or dual controller
- Two 1-GbE ports per controller
- Multiple RAID levels (0, 1, 3, 5, 6, 10, and 50)
- MSA2000 drive enclosures (MSA70 supported with MSA2000i G2)
- Path failover capability

For more information, see the *HP StorageWorks 2312i Modular Smart Array User Guide*.

### P2000 G3 FC/iSCSI, P2000 G3 10Gb iSCSI, P2000 G3 iSCSI overview

HP StorageWorks P2000 G3 family of iSCSI products are controller shelves with a network interface to serve as an iSCSI target. The three different controller types give exceptional flexibility for a wide range of customer installations.

Table 188 (page 318) describes the P2000 G3 family controller types.

**Table 188 P2000 G3 controller types**

Controller	FW version	Redundancy	Interface	Expansion
P2000 G3 FC/iSCSI	TS201R0xx	Single or dual controller	Two 1-GbE iSCSI <sup>1</sup>	P2000 LFF drive enclosure D2700 SFF disk enclosure
P2000 G3 10-GbE iSCSI	TS210R0xx		Two 10-GbE SFP+ <sup>2</sup>	
P2000 G3 iSCSI	TS220R0xx		Four 1-GbE iSCSI	MSA2000 LFF drive enclosure MSA70 SFF drive enclosure

<sup>1</sup> Controller also includes two 8Gb Fibre Channel ports described in “P2000 and MSA storage system rules” (page 203).

<sup>2</sup> SFP+ optical transceivers or SFP+ direct attached cables must be purchased separately.

## P2000 G3 FC/iSCSI, MSA2000i G2 and MSA2000i support

Table 189 (page 319) describes the P2000 G3 FC/iSCSI, MSA2000i G2 and MSA2000i support iSCSI initiators for Microsoft Windows and Linux.

**Table 189 P2000 G3 FC/iSCSI, MSA2000i G2, and MSA2000i configurations**

Model	Array chassis	Expansion <sup>1</sup>	Expansion maximum
P2000 FC/iSCSI P2000 G3 10-GbE iSCSI P2000 G3 iSCSI	P2000 3.5-in drive bay (LFF) P2000 2.5-in drive bay (SFF) 2012 3.5-in drive bay (LFF, upgrade only) 2024 2.5-in drive bay (SFF, upgrade only)	P2000 LFF drive enclosure D2700 SFF disk enclosure MSA2000 LFF drive enclosure MSA70 SFF drive enclosure	8 total enclosures or 149 drives 96 LFF drives 149 SFF drives
MSA2000i G2 MSA2300i	2012 3.5-in drive bay (LFF) 2024 2.5-in drive bay (SFF)	MSA2000 LFF drive enclosure MSA70 SFF drive enclosure	5 total enclosures or 99 drives 60 LFF drives 99 drives
MSA2012i	Pre-configured single or dual controller 3.5-in drive bay (LFF)	MSA2000 LFF drive enclosure	4 total enclosures 48 LFF drives

<sup>1</sup> Confirm cabling requirements and limitation in the Cable Configuration Guide at <http://www.hp.com/go/msa>

### Server support

The MSA2000i G2 and MSA2000i support the following HP ProLiant servers:

- HP ProLiant DL, ML
- HP ProLiant c-Class BladeSystem

---

**NOTE:** BL20p G1 servers are not supported.

---

## Operating system support

Table 190 (page 320) lists the supported operating systems for the MSA2000i G2 and MSA2000i.

**Table 190 MSA2000i G2 and MSA2000i operating system support**

Storage system	Operating systems
P2000 G3 FC/iSCSI P2000 G3 10GbE iSCSI P2000 G3 iSCSI	Microsoft Windows Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX
MSA2000i G2 MSA2300i	Microsoft Windows Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX Sun Solaris Citrix XenServer
MSA2000i	Microsoft Windows Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX

## Path failover software

The P2000 G3 FC/iSCSI, MSA2000i G2 and MSA2000i support the following:

- Microsoft MPIO basic failover software for Windows operating systems
- Device Mapper for path failover with Linux operating systems

## Management software support

The P2000 G3 family, MSA2000i G2 and MSA2000i support target-based management interfaces, including Telnet (CLI), FTP, and a web-based interface. The web-based interface is supported with Microsoft Internet Explorer 6 (or later) and Mozilla Firefox 1.0.7 (or later).



## Maximum configurations

Table 191 (page 321) lists the maximum configurations for the MSA2000i G2 and MSA2000i.

**Table 191 MSA2000i G2 and MSA2000i maximum configurations**

Storage systems	Operating systems	Drives	Hosts	LUNs	LUN size	Snapshots and volume copies <sup>1</sup>
P2000 G3 FC/iSCSI P2000 G3 10GbE iSCSI	Microsoft Windows Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX	149 SFF 96 LFF	64	512	16 TB	Up to 64 snapshots Additional license to 512
MSA2000i G2 MSA2324i G2 MSA2312i G2	Microsoft Windows Red Hat Enterprise Linux SUSE Linux Enterprise Server VMware ESX Sun Solaris	99 SFF 60 LFF	64	512	16 TB	Up to 256 snapshots Up to 256 volume copies
MSA2000i	Microsoft Windows Red Hat Enterprise Linux VMware ESX	48	16	256 <sup>2</sup>	16 TB	Up to 64 snapshots Up to 128 volume copies

<sup>1</sup> Snapshots and volume copies require additional licenses.

<sup>2</sup> A single controller supports 128 LUNs. Two controllers are required for 256 LUNs.

## HP StorageWorks P4000 SAN

This section describes the following topics:

- “P4000 SAN Solutions overview” (page 321)
- “P4000 SAN Solutions support” (page 322)

In this section, “P4000 SAN” refers to both the P4000 SAN and P4000 G2 SAN products.

### P4000 SAN Solutions overview

HP StorageWorks P4000 SANs are composed of multiple storage nodes consolidated into single or multiple pools of storage resources. All available capacity and performance is aggregated and available to every volume in the cluster. With multiple iSCSI network interfaces across the cluster, a Virtual IP address across these interfaces serves as an iSCSI target.

Key features include:

- A cluster of pooled storage supporting up to 16 storage nodes. For larger configurations, please contact an HP storage representative.
- Two 1-GbE ports per storage node (or optional upgrade to SPF+ 10-GbE dual ports per storage node).
- Multiple layers of high availability:
  - Multiple RAID levels within the storage node (5, 6, and 10)
  - Multiple Network RAID levels across the storage nodes on a per volume basis (RAID 0, 10, 10+1, 10+2, 5, 6)
- P4300 SAN nodes support 8 disk drives. P4500 SAN nodes support 12 disk drives. P4800 SAN nodes support 35 disk drives.

- Path failover capability:
  - Within the storage node, iSCSI interface bonding (active/passive, ALB, 802.3ad LACP).
  - Within the cluster, Network RAID configured volumes support another path to another storage node.
- All P4000 SANs come with an all-inclusive enterprise feature set which includes: storage clustering, network RAID, thin provisioning.

For more information and additional resources, including links to manuals, see [www.hp.com/go/P4000Support](http://www.hp.com/go/P4000Support) and *Building High-performance, High availability IP Storage Networks with HP SAN/iQ Software*.

## P4000 SAN Solutions support

For the latest information on version support, see the *HP StorageWorks P4000 SAN Solutions Compatibility Matrix* at <http://www.hp.com/go/P4000Compatibility> or on SPOCK at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access to SPOCK.

### Path failover software

The P4000s support the following:

- Microsoft Native MPIO, and the HP 4000 DSM for MPIO with Windows server operating systems
- VMware Round Robin MPIO
- Citrix/Linux/Unix/VMware bonding of network interfaces performed at the networking layer
- For more information, see the following documents, available at [www.hp.com/go/P4000Support](http://www.hp.com/go/P4000Support) (see Manuals section):
  - *HP LeftHand P4000 Windows Solution Pack User Guide*
  - *Running VMware vSphere 4 on HP LeftHand P4000 SAN Solutions*

### Management software support

The P4000s support target-based management interfaces, including the HP P4000 Centralized Management Console (CMC), Command Line Interface (CLIQ), and SSH.

For more information, see the following documents, available at [www.hp.com/go/P4000Support](http://www.hp.com/go/P4000Support) (see Manuals section):

- *HP StorageWorks P4000 SAN/iQ Command-Line Interface User Manual*
- *HP StorageWorks SAN Solutions Storage System User Guide*

### Maximum configurations

The P4000 SAN Solutions Configuration Summary reports information about storage items, including color-coded status regarding recommended limits based on performance and scalability.

- When a configuration category is nearing the maximum recommended limit, the navigation window displays the category information as orange.
- When a configuration category reaches the maximum recommended limit, the navigation window displays the category information as red.
- When the number in that category is reduced, the color changes immediately to reflect the new state.

For example, if you have numerous schedules for a large number of volumes that are creating and deleting snapshots. When the number of snapshots approaches the maximum recommended number, the summary bar changes from green to orange. After enough snapshots are deleted from the schedule, the summary bar returns to green.

## Best practices

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

Table 192 (page 323) describes the recommended maximum configurations.

**Table 192 Configuration recommendations**

Storage item	Best practice	Caution	Not recommended
Status indicator	Green	Orange	Red
Volumes + Snapshots + SmartClones	1 to 1,000	1,001 to 1,500	1,501+
iSCSI sessions per group	1 to 4,000	4,001 to 5,000	5,001+
Nodes per group	1 to 20	21 to 32	33+
Nodes per cluster	1 to 10	11 to 15	16+

## LUN size

The maximum LUN size for a fully provisioned volume depends on the storage node size, local RAID configuration, and network RAID configuration settings for the volume. Thinly-provisioned volumes can physically grow only up to the maximum fully-provisioned LUN equivalent in the cluster.

## HP iSCSI bridge products

### Bridging and routing

iSCSI routers and bridges are gateway devices that connect storage protocols such as Fibre Channel or SCSI to IP networks. iSCSI routers and bridges enable block-level access across networks. Routing data requests from an IP network device to a Fibre Channel device involves these steps:

1. An iSCSI host makes a storage data request.
2. The request is switched (or routed) through the IP network with the destination IP address of an iSCSI bridge or router.
3. The bridge or router converts the iSCSI request into its Fibre Channel equivalent.
4. The converted request is sent to the Fibre Channel target storage device.

The bridge or router performs the reverse conversion as the Fibre Channel target responds to the iSCSI host. Conversion is transparent to both host and target since the iSCSI bridge or router mediates the exchange at wire speed.

### iSCSI bridge to Fibre Channel

Many iSCSI bridges are compatible with HP-supported Fibre Channel switches. This bridging technology routes Fibre Channel storage to all servers in an IP fabric.

IP hosts use iSCSI to access Fibre Channel storage systems. The storage systems appear as direct-attached storage to the hosts.

Key benefits of bridging include:

- Consolidating Fibre Channel storage to both Fibre Channel and IP hosts
- Extending Fibre Channel distance limitations through access to IP networks

The HP-supported iSCSI to Fibre Channel bridge products are as follows:

- B-series iSCSI Director Blade
- B-series MP Router
- C-series IP Storage Services Modules
- C-series 14/2 Multiprotocol Services Module
- C-series 18/4 Multiprotocol Module
- MPX200 Multifunction Router iSCSI
- EVA iSCSI Connectivity Option
- EVA4400 iSCSI Connectivity Option
- SR2122-2 IP Storage Router

## MPX200 Multifunction Router with iSCSI for EVA storage

The P6000/EVA family of FC storage systems supports integrated iSCSI connectivity using the MPX200 Multifunction Router. The MPX200 hardware is integrated with up to four P6000/EVA or XP 24000/20000 storage systems (see [“MPX200 Multifunction Router with iSCSI for XP storage” \(page 330\)](#)) and HP P6000 Command View to deliver multiprotocol capabilities. This provides iSCSI and FC attached servers access to block storage through an FC and Ethernet IP network simultaneously. The MPX200 is available from HP factory-integrated with a P6000/EVA storage system or as a field upgrade to an existing P6000/EVA storage system. With this product, iSCSI connectivity to the P6000/EVA is provided for servers through a standard 1-GbE or 10-GbE NIC.

The MPX200 chassis contains one or two router blades, two PCMs, and a midplane. There are two types of router blades: a 4-port 1-GbE blade and a 2-port 10-GbE/2-port 1-GbE blade. Both blade options include two 8-Gb/s FC ports. MPX200 dual-blade configurations provide for high availability with failover between the blades.

### MPX200 simultaneous operation

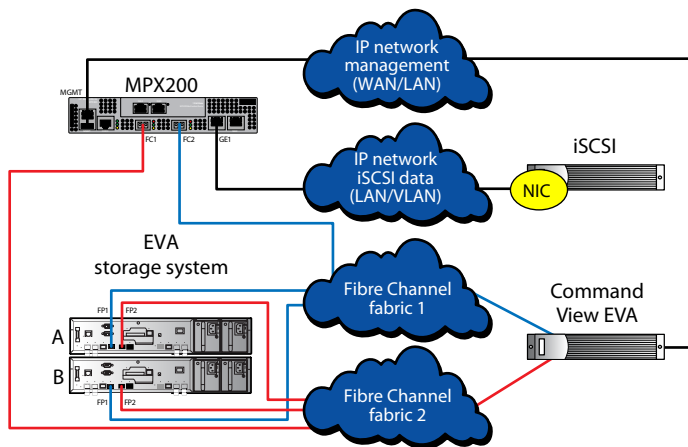
The MPX200 Multifunction Router supports iSCSI, FCoE, data migration, and FCIP. The base functionality is iSCSI/FCoE, with the option to add either data migration or FCIP for standalone or concurrent operation. This section describes iSCSI usage and support. For information about using the MPX200 FCoE feature, see [“Fibre Channel over Ethernet” \(page 61\)](#). For information about using the MPX200 data migration feature, see [“MPX200 Multifunction Router with data migration” \(page 218\)](#). For information about using the MPX200 FCIP feature, see [“MPX200 Multifunction Router with FCIP” \(page 276\)](#). For information about configuring the MPX200 for multiple functions concurrently, see the *HP StorageWorks MPX200 Multifunction Router User Guide*.

### MPX200 configuration options

A P6000/EVA storage system can be configured for simultaneous connectivity to iSCSI and FC attached hosts. Support for iSCSI to a P6000/EVA storage system is provided through the MPX200 and an existing FC switch fabric port (fabric-attached) or direct-connect to an EVA controller port.

[Figure 133 \(page 325\)](#) shows an MPX200-EVA single-blade fabric-attached configuration. This is the lowest-cost configuration and is used when high availability for iSCSI hosts is not required.


**Figure 133 MPX200-EVA single-blade fabric-attached configuration**



26534a

Figure 134 (page 325) shows an MPX200-EVA dual-blade fabric-attached configuration. This configuration provides high availability with failover between blades.

**Figure 134 MPX200-EVA dual-blade fabric-attached configuration**



26550a

Figure 135 (page 326) shows a multi-EVA configuration with connectivity for up to four EVA storage systems from a single MPX200 blade.

Figure 135 MPX200 single-blade multi-EVA configuration

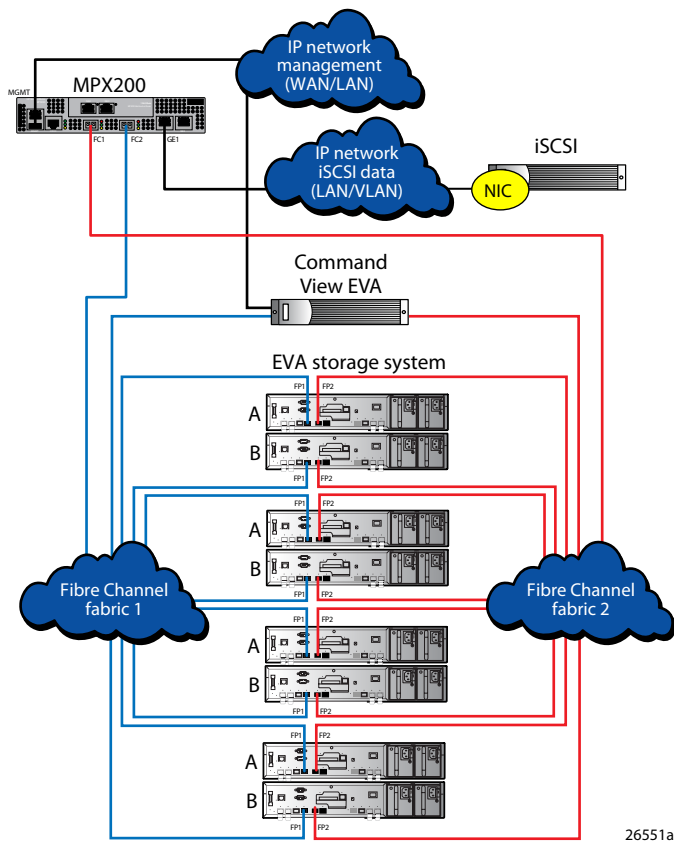


Figure 136 (page 327) shows a multi-EVA configuration with connectivity for up to four EVA storage systems from dual MPX200 blades. This configuration provides high availability with failover between blades.

Figure 136 MPX200 dual-blade multi-EVA configuration

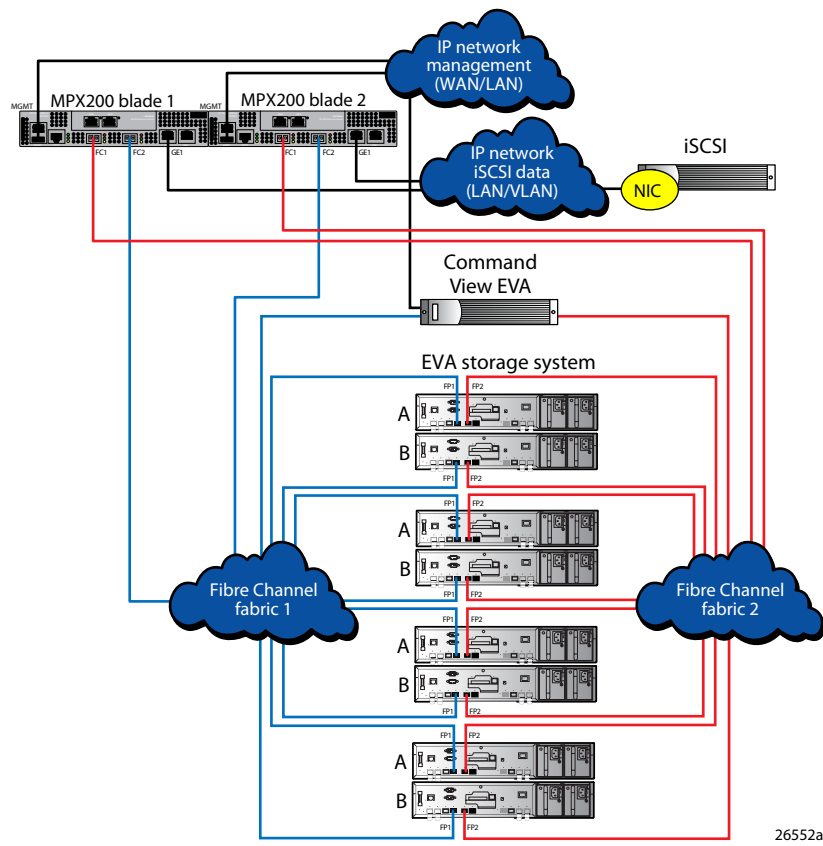
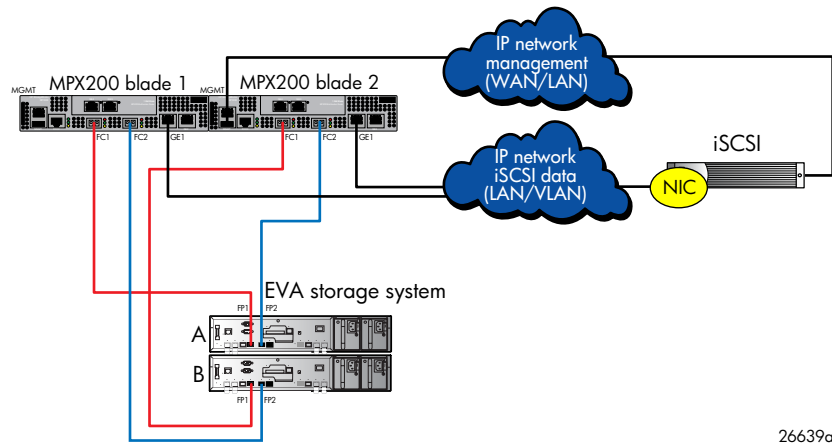
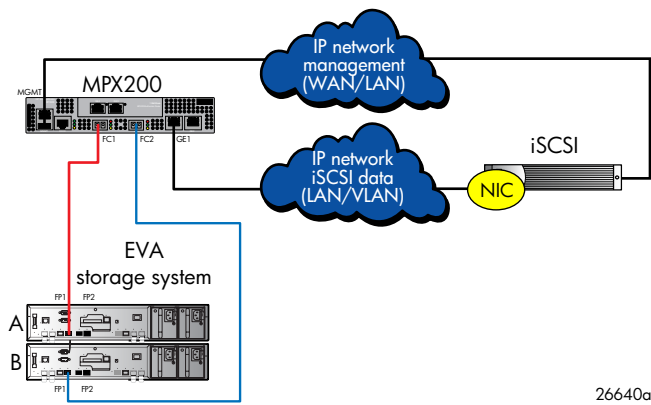


Figure 137 (page 327), Figure 138 (page 328), and Figure 139 (page 328) illustrate EVA direct connect configurations.

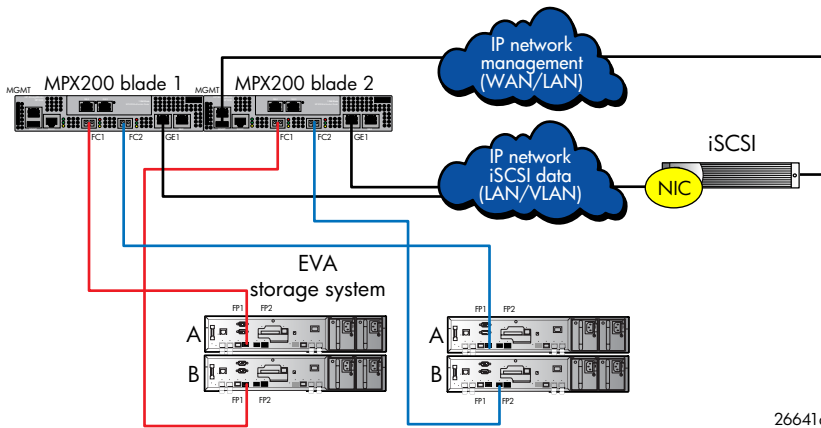
Figure 137 MPX200 dual-blade direct connect to one EVA configuration



**Figure 138 MPX200 single-blade direct connect to one EVA configuration**



**Figure 139 MPX200 dual-blade direct connect to two EVA configuration**



## MPX200 iSCSI rules and supported maximums

The MPX200 chassis can be configured with one or two blades. Dual-blade configurations provide for high availability with failover between blades. All dual-blade configurations are supported as redundant pairs only. iSCSI-connected servers can be configured for access to one or both blades.

**NOTE:** In the event of a failover between blades, servers with single-blade connectivity to a failed blade will no longer have connectivity to the MPX200.

Table 193 (page 328) lists the supported maximums for the MPX200.

**Table 193 Supported MPX200 maximums**

Description	Maximum per MPX200 solution <sup>1</sup>
<b>Hardware</b>	
P6000/EVA and/or XP24000/20000 storage systems (see “MPX200 Multifunction Router with iSCSI for XP storage” (page 330))	4 total (any combination)
MPX200	1 chassis with up to 2 blades
MPX200 iSCSI port connections	See the <i>HP StorageWorks MPX200 Multifunction Router User Guide</i> at <a href="http://hp.com/go/mpx200">http://hp.com/go/mpx200</a> .
<b>Configuration parameter</b>	
Total number of iSCSI initiators	300 per chassis for 1-GbE (1 or 2 blades) 600 per chassis for 10-GbE (1 or 2 blades)



**Table 193 Supported MPX200 maximums (continued)**

Description	Maximum per MPX200 solution <sup>1</sup>
Total number of iSCSI LUNs	4,096 per chassis, 1,024 per EVA or XP
iSCSI connections, 1-GbE	1,024 per blade, 2,048 per chassis
iSCSI connections, 10-GbE	2,048 per blade, 4,096 per chassis

<sup>1</sup> For mixed blade type chassis configurations that include one 1-GbE blade and one 10-GbE blade, the maximums values for a 1-GbE blade prevail.

## P6000/EVA storage system rules and guidelines

The MPX200 is supported with the following P6000/EVA storage systems:

- EVA4400/4400 with embedded switch
- EVA4000/4100/6000/6100/8000/8100
- EVA6400/8400
- P6300/P6500

All MPX200 configurations must follow these P6000/EVA connectivity rules:

- When using the MPX200 for iSCSI, MPX200 FC connections can be direct connect to a P6000/EVA controller host port or fabric connect through an FC switch.
- Each P6000/EVA storage system can connect to a maximum of one MPX200 chassis (two blades).
- Each P6000/EVA controller host port can connect to a maximum of two MPX200 FC ports.
- A maximum of one MPX200 chassis (two blades) can be zoned with up to four P6000/EVA storage systems.
- A P6000/EVA storage system can present LUNs to iSCSI initiators and FC hosts concurrently.

Table 194 (page 329) describes the maximum number of servers, initiators, and LUNs supported for combined iSCSI and Fibre Channel connectivity to a P6000/EVA.

**Table 194 Supported P6000/EVA/MPX200 maximums**

	P6000/EVA with Fibre Channel only	P6000/EVA with Fibre Channel and 1-GbE iSCSI	P6000/EVA with Fibre Channel and 10-GbE iSCSI
Maximum number of servers	256	552 with 1 EVA 1,320 with 4 EVAs	852 with 1 EVA 1,608 with 4 EVAs
Maximum number of initiators	1,024	1,308 with 1 EVA 4,332 with 4 EVAs	1,608 with 1 EVA 4,632 with 4 EVAs
Maximum number of LUNs <sup>1</sup> 4,096 total per MPX200 chassis	1,023 EVA4x00 /6000/6100 /8000/8100 2,047 EVA6400/8400	1,023 EVA4x00/6000/6100/8000/8100 2,047 EVA6400/8400	

<sup>1</sup> For more information, see "Configuration parameters" (page 216).

## HP P6000 Command View and MPX200 management rules and guidelines

The HP P6000 Command View implementation for the MPX200 supports management of up to four EVA storage systems concurrently. This implementation provides the equivalent functionality for both iSCSI and FC connected servers. All MPX200 management functions are integrated in HP P6000 Command View. For more information, see the *HP StorageWorks MPX200 Multifunction Router User Guide* at <http://hp.com/go/mpx200>.

Observe the following MPX200 HP P6000 Command View rules and guidelines:

- Command View EVA 9.1.1 (or later) for server based or array based management.
- A maximum of one MPX200 chassis (two blades) can be discovered by an EVA storage system.
- HP P6000 Command View manages the MPX200 out of band (IP) through the MPX200 management IP port. The HP P6000 Command View application server must be on the same IP network as the MPX200 management IP port.
- The MPX200 iSCSI initiator or iSCSI LUN masking information does not reside in the HP P6000 Command View database. All iSCSI initiator and LUN presentation information resides in the MPX200.
- The default iSCSI initiator EVA host mode setting is Microsoft Windows. The iSCSI initiator for Apple Mac OS X, Linux, Sun Solaris, VMware, and Windows 2008 host mode setting is configured with HP P6000 Command View.

---

**NOTE:** Communication between HP P6000 Command View and the MPX200 is not secured by the communication protocol. If this unsecured communication is a concern, HP recommends a confined or secured IP network within a data center.

---

### EVA storage system software

For FCIP, the MPX200 is supported with HP P6000 Continuous Access, Business Copy, SSSU, or Replication Solutions Manager.

### Fibre Channel switch and fabric support

The MPX200 is supported with HP B-series, C-series, and H-series switch models. For the latest information on FC switch model and firmware support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

### Operating system and multipath software support

This section describes the MPX200 iSCSI operating system, multipath, and cluster support.

For the latest information on operating system and multipath software support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 195 (page 330) lists operating system and multipath software support.

**Table 195 MPX200-EVA operating system, multipath software, and cluster support**

Operating system	Multipath software	Clusters	P6000/EVA storage system
Apple Mac OS X	None	None	EVA4400/4400 with embedded switch EVA4000/4100/6000/6100/8000/8100 EVA6400/8400 P6300/P6500
Microsoft Windows Server 2008, 2003 Hyper-V	MPIO with HP DSM MPIO with Microsoft DSM	MSCS <sup>1</sup>	
Linux Red Hat, SUSE	Device Mapper	None	
Sun Solaris	Solaris MPxIO	None	
VMware	VMware MPxIO	None	

<sup>1</sup> MSCS is supported with Microsoft DSM only.

### MPX200 Multifunction Router with iSCSI for XP storage

The XP24000/20000 family of FC storage systems supports integrated iSCSI connectivity using the MPX200 Multifunction Router. The MPX200 hardware is integrated with up to four XP or EVA storage systems (see “MPX200 Multifunction Router with iSCSI for EVA storage” (page 324)) to

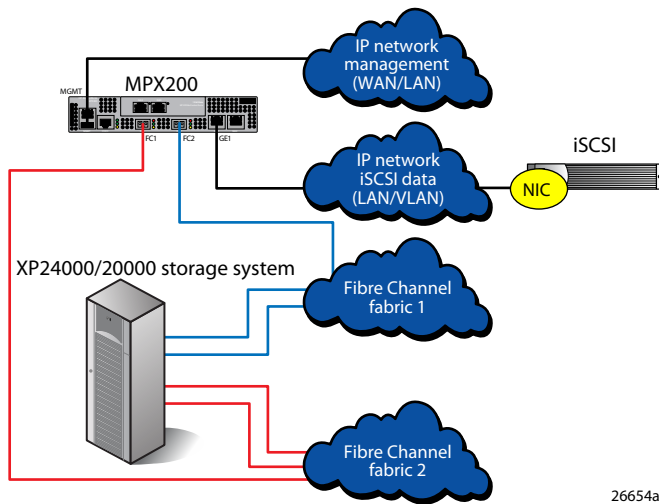
deliver multiprotocol capabilities. This provides iSCSI and FC attached servers access to block storage through FC and Ethernet IP network simultaneously. The MPX200 is available from HP as an option for a new XP24000/20000 purchase or as a field upgrade to an existing XP24000/20000 storage systems. With this product, iSCSI connectivity to the XP is provided for servers through a standard 1-GbE or 10-GbE NIC.

### MPX200 configuration options

An XP storage system can be configured for simultaneous connectivity to iSCSI and FC attached hosts. Support for iSCSI to an XP storage system is provided through the MPX200 and an existing FC switch fabric port (fabric-attached).

Figure 140 (page 331) shows an MPX200-XP single-blade fabric-attached configuration. This is the lowest-cost configuration and is used when high availability for iSCSI hosts is not required.


**Figure 140 MPX200-XP single-blade fabric-attached configuration**



26654a

Figure 141 (page 331) shows an MPX200-XP dual-blade fabric-attached configuration. This configuration provides high availability with failover between blades.

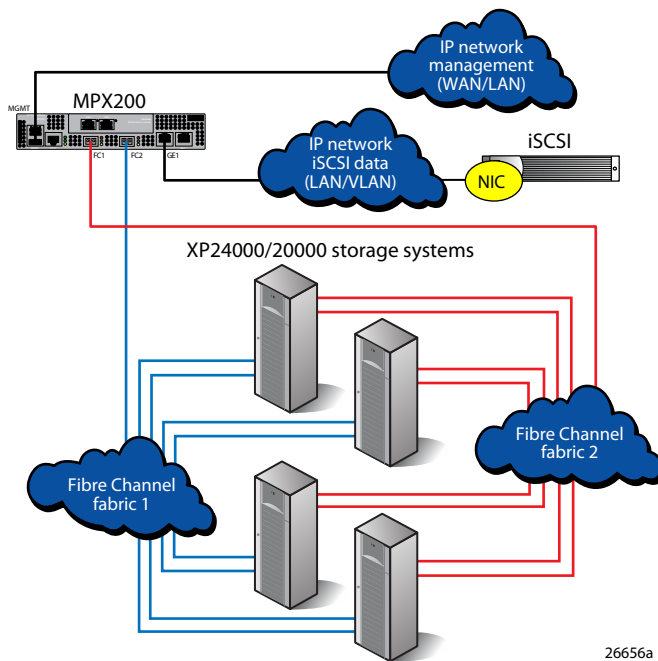
**Figure 141 MPX200-XP dual-blade fabric-attached configuration**



26655a

Figure 142 (page 332) shows a multi-XP configuration with connectivity for up to four XP storage systems from a single MPX200 blade.

**Figure 142 MPX200-XP multi-XP fabric-attached configuration**



## MPX200 iSCSI rules and supported maximums

The MPX200 chassis can be configured with one or two blades. Dual-blade configurations provide for high availability with failover between blades. All dual-blade configurations are supported as redundant pairs only. iSCSI-connected servers can be configured for access to one or both blades.

**NOTE:** In the event of a failover between blades, servers with single-blade connectivity to a failed blade will no longer have connectivity to the MPX200.

Table 193 (page 328) lists the supported iSCSI connectivity maximums for the MPX200. For information about XP24000/20000 FC host connectivity, see the XP24000/20000 documentation.

## XP storage system rules and guidelines

The MPX200 is supported with the following XP storage systems:

- XP24000
- XP20000

All MPX200 configurations must follow these connectivity rules:

- When using the MPX200 for iSCSI, MPX200 FC connections must be fabric-attached through an FC switch
- Each XP storage system can connect to a maximum of one MPX200 chassis (two blades)
- A maximum of one MPX200 chassis (two blades) can be zoned with up to four XP storage systems
- XP and EVA storage systems can connect to the same MPX200, the total allowable number of storage systems is four per MPX200 chassis
- An XP storage system can present LUNs to iSCSI initiators and FC hosts concurrently

## Operating system and multipath software support

This section describes the MPX200 iSCSI operating system and multipath support.

For the latest information on operating system and multipath software support, see the SPOCK website at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Table 196 (page 333) lists operating system and multipath software support.

**Table 196 MPX200-XP operating system and multipath support**

Operating system	Multipath software	XP storage system
Microsoft Windows Server 2008, 2003	MPIO with Microsoft DSM	XP24000/20000

## EVA and EVA4400 iSCSI Connectivity Option

The EVA iSCSI Connectivity Option allows iSCSI connectivity support for the EVA family of storage systems. The EVA iSCSI Connectivity Option uses HP P6000 Command View management software and the following hardware:

- mpx100 for all EVA models
- mpx100b for EVA4400 and EVA4400 (embedded switch)

The EVA iSCSI Connectivity Option uses a standard GbE NIC for iSCSI connectivity to servers. You can purchase this option either included with an EVA or as a field upgrade to your existing EVA.

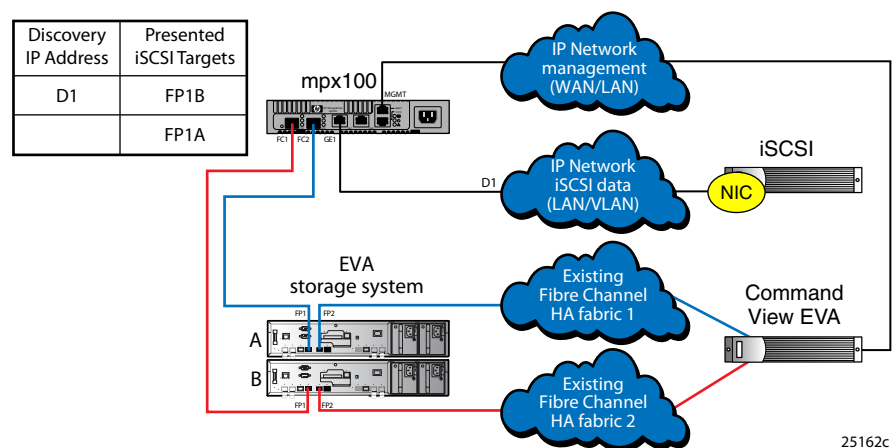
**NOTE:** The EVA iSCSI Connectivity Option (mpx100) is supported with EVA4000/4100/4400/4400 (embedded switch)/6000/6100/6400/8000/8100/8400 and EVA3000/5000 storage systems. The EVA4400 iSCSI Connectivity Option (mpx100b) is supported with EVA4400 and EVA4400 with the embedded switch storage systems.

For additional product information, including product documentation, see [http://h18006.www1.hp.com/products/storageworks/evaiscsconnect/index.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/products/storageworks/evaiscsconnect/index.html?jumpid=reg_R1002_USEN).

An EVA storage system can connect simultaneously to iSCSI and Fibre Channel attached hosts. iSCSI support is provided through a dedicated EVA host port for direct connect (see Figure 143 (page 333), Figure 144 (page 334), and Figure 146 (page 335)); or shared with Fibre Channel through an existing fabric host port for fabric attachment (see Figure 147 (page 335)).

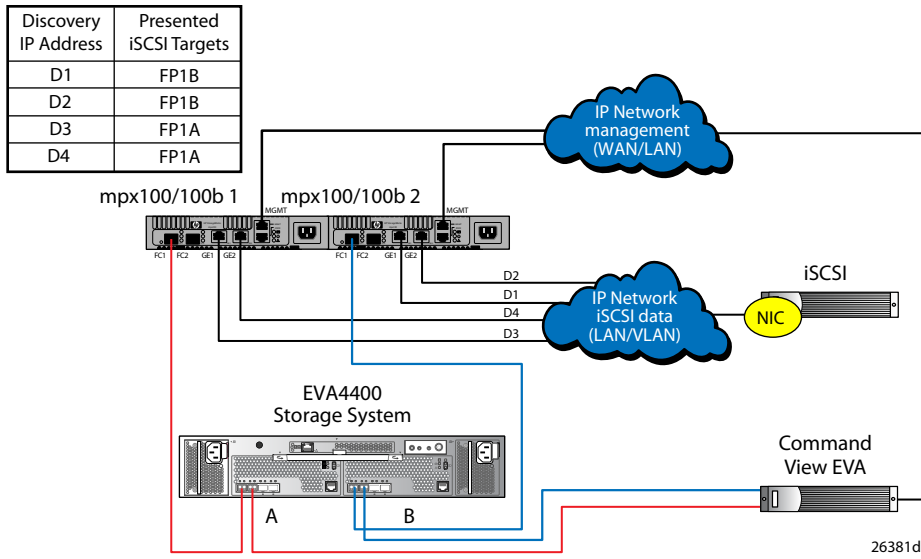
The EVA4400 with the ABM module can be configured with direct connections as shown in Figure 145 (page 334) as an iSCSI-only solution.

**Figure 143 Direct connect iSCSI-Fibre Channel attachment mode configuration**

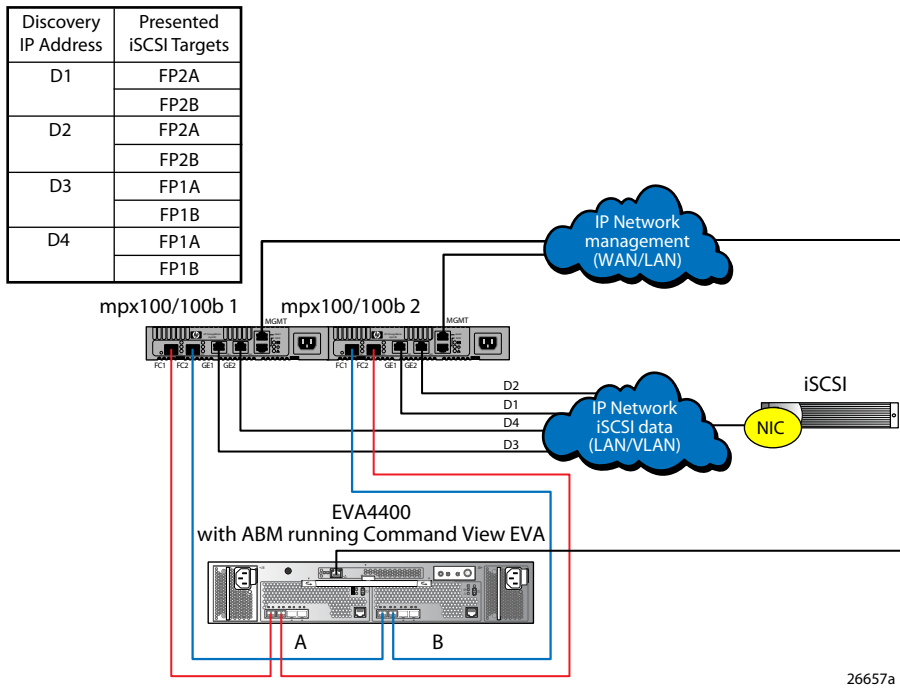


**NOTE:** Direct connect mode requires a dedicated host port on each HSV controller.

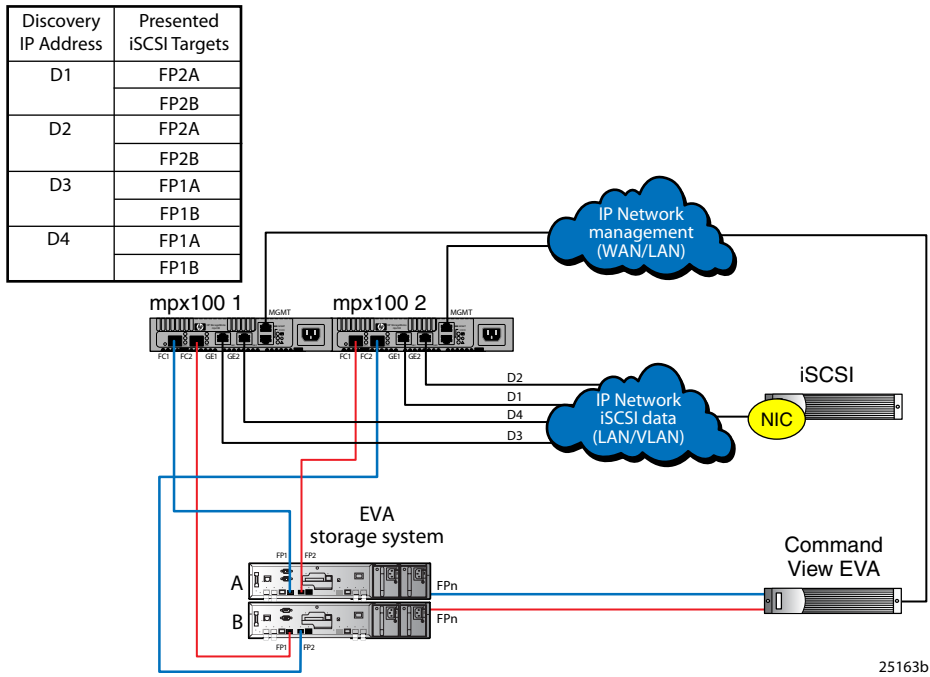
**Figure 144 EVA4400 direct connect iSCSI-Fibre Channel attachment mode configuration**



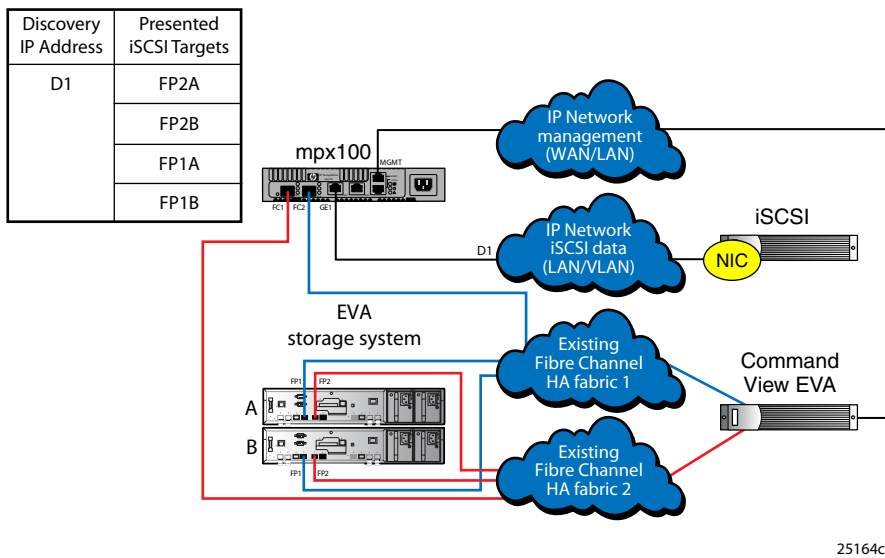
**Figure 145 EVA4400 direct connect iSCSI-Fibre Channel attachment mode ABM configuration**



**Figure 146 EVA8000/8100 mpx100 and Windows host direct connect configuration**



**Figure 147 Fabric iSCSI-Fibre Channel attachment mode configuration**



## Hardware support

This section describes the hardware devices supported by the EVA and EVA4400 iSCSI Connectivity Options.

### mpx100/100b data transport

The EVA and EVA4400 iSCSI options support both direct connect and Fibre Channel fabric connectivity through the mpx100/100b to the EVA storage system.

Table 197 (page 336) describes the connectivity attachment mode based on the EVA storage system model.

**Table 197 EVA storage system connectivity attachment modes**

EVA storage system	Storage software version	iSCSI-Fibre Channel attachment mode <sup>1</sup>
EVA4400	For the latest information on storage software version support, see the product release notes or SPOCK at <a href="http://www.hp.com/storage/spock">http://www.hp.com/storage/spock</a> . You must sign up for an HP Passport to enable access.	mpx100b direct connect (Figure 143 (page 333))
EVA4400 Embedded Switch Module, 8 Gb Brocade		mpx100b fabric through a Fibre Channel switch (Figure 147 (page 335))
EVA4000/4100/6000/6100/6400/8000/8100/8400		mpx100b fabric through the embedded Fibre Channel switch
EVA3000/5000		mpx100 direct connect (Figure 143 (page 333))
		mpx100 and Windows host direct connect only (all controller host ports direct connect) (Figure 146 (page 335))
		mpx100 fabric through a Fibre Channel switch (Figure 147 (page 335))
		mpx100 fabric through a Fibre Channel switch (Figure 147 (page 335))

<sup>1</sup> A Fibre Channel switch is not required for mpx100 and Windows host direct connect only or HP P6000 Command View iSCSI deployment. For more information, see Figure 143 (page 333), Figure 144 (page 334), and Figure 146 (page 335).

### Fibre Channel switches

The EVA and EVA4400 iSCSI Connectivity Options are supported with most B-series, C-series, and M-series switches. For Fibre Channel switch support, see:

- “B-series switches and fabric rules” (page 83)
- “C-series switches and fabric rules” (page 115)
- “M-series switches and fabric rules” (page 140)

**NOTE:** Not all switch models are supported with the EVA iSCSI Connectivity Options. Contact an HP storage representative for the latest information about switch model support.

### Storage systems

The mpx100 supports EVA4000/4100/4400/4400 (embedded switch)/6000/6100/6400/8000/8100/8400 and EVA3000/5000 storage systems. The mpx100b supports the EVA4400 storage system only.

### Software support

This section describes the software supported by the EVA iSCSI Connectivity Option.

#### Management software

The required minimum versions of HP P6000 Command View/Command View EVA management software are as follows:

- HP P6000 Command View 9.4 for P6300/P9500 EVA
- Command View 9.0.1 (or later) is required for EVA6400/8400.
- Command View 8.1 (or later) is required for array-based management.
- Command View 8.0.1 (or later) is required for EVA4400 (embedded switch) storage systems.

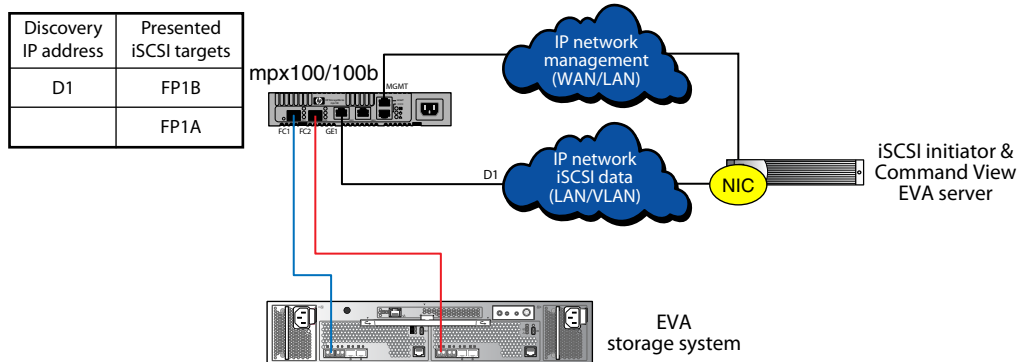


- Command View 8.0x (or later) is required for the mpx100/100b running firmware version 2.4x (or later).
- The EVA iSCSI Connectivity Options support Command View EVA iSCSI connectivity (Fibre Channel switch not required).

See [Figure 148 \(page 337\)](#) and [Figure 149 \(page 337\)](#).

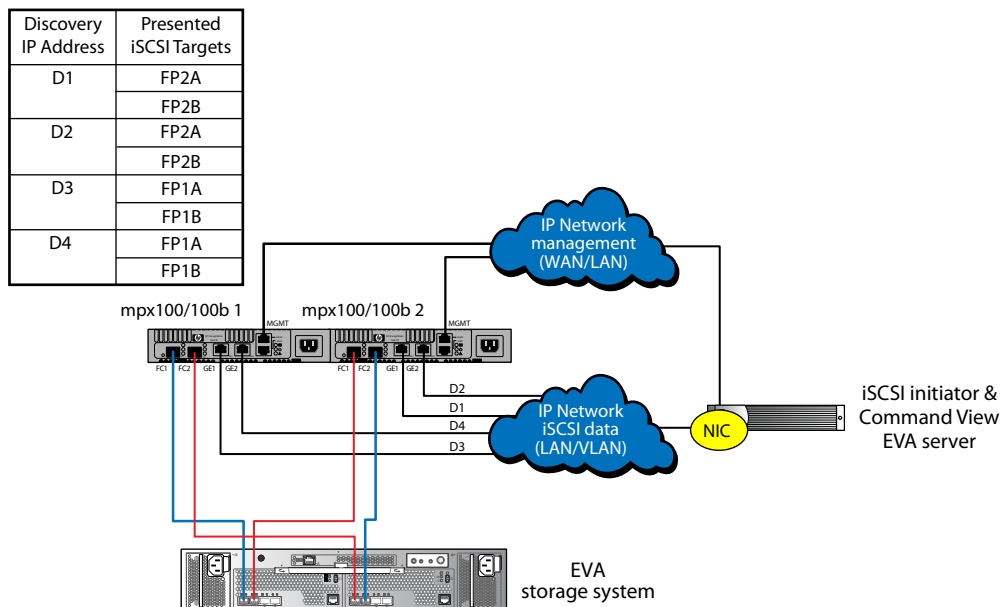
**NOTE:** HP StorageWorks mpx Manager is required for mpx100/100b management.

**Figure 148 Command View EVA iSCSI connectivity configuration 1**



26484b

**Figure 149 Command View EVA iSCSI connectivity configuration 2**



26485b

### Multipath software

The EVA iSCSI Connectivity Option supports iSCSI multipath connectivity on HP OpenVMS, Microsoft Windows with MPIO, Linux with Device Mapper, and Sun Solaris and VMware ESX with MPxIO. iSCSI multipath connectivity is supported on EVA4000/4100/4400/4400 (embedded switch)/6000/6100/6400/8000/8100/8400 and the EVA3000/5000 (active/active).

The EVA4400 iSCSI Connectivity Option supports iSCSI multipath connectivity on Microsoft Windows with MPIO, Linux with Device Mapper, and Sun Solaris and VMware ESX with MPxIO. iSCSI multipath connectivity is supported on EVA4400 using XCS 09000000 (minimum) or EVA4400 (embedded switch) using XCS 09003000 (minimum).

**NOTE:** For the latest information on version support, see the product release notes or SPOCK at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

Figure 150 (page 338) illustrates the high-availability, multipath direct connect iSCSI-Fibre Channel attachment mode configuration with Windows MPIO.

**Figure 150 Multipath direct connect iSCSI-Fibre Channel attachment mode configuration**

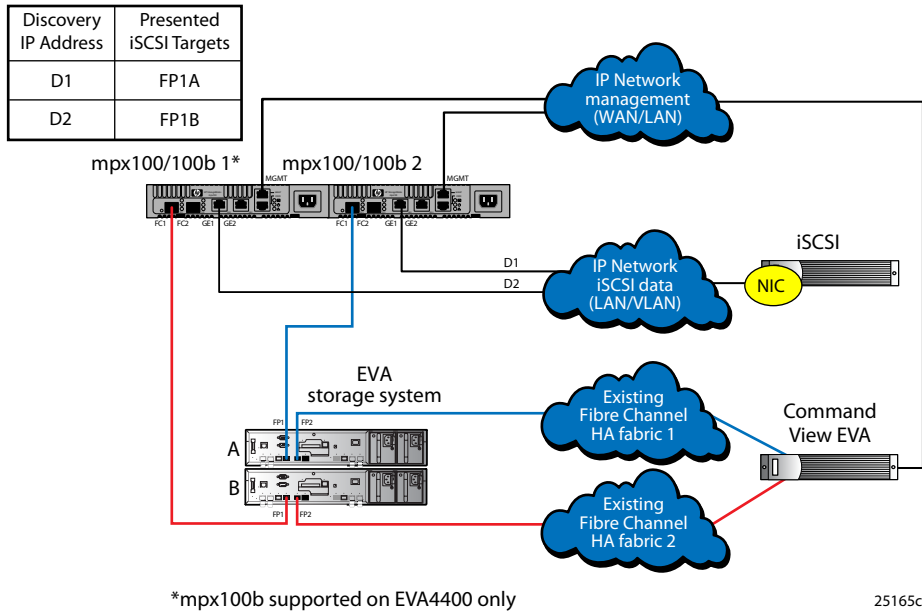
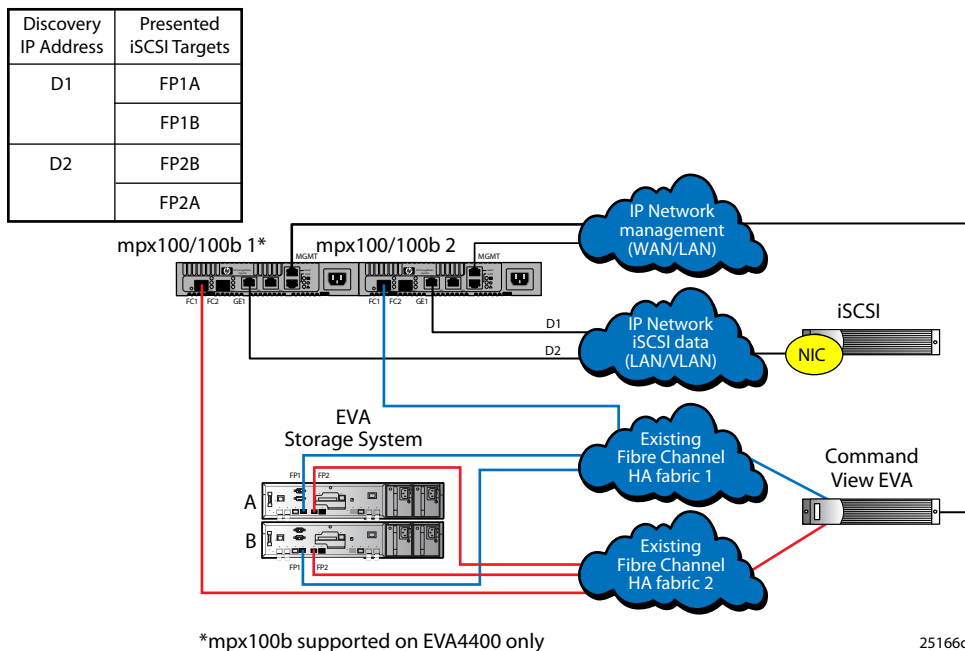


Figure 151 (page 338) illustrates the high-availability multipath fabric iSCSI-Fibre Channel attachment mode configuration.

**Figure 151 Multipath fabric iSCSI-Fibre Channel attachment mode configuration**



### HP P6000 Continuous Access

An EVA LUN that has been presented to an iSCSI initiator is supported with current EVA storage software applications such as HP P6000 Continuous Access, Business Copy, SSSU, and Replication

Solutions Manager. See the *EVA iSCSI Connectivity Option User Guide* at [http://h18006.www1.hp.com/products/storageworks/evaiscsiconnect/index.html?jumpid=reg\\_R1002\\_USEN](http://h18006.www1.hp.com/products/storageworks/evaiscsiconnect/index.html?jumpid=reg_R1002_USEN).

### Operating systems and network interface cards

The EVA and EVA4400 iSCSI options support the following operating systems, unless noted otherwise:

- Apple Mac OS X
- Linux—Red Hat
- Linux—SUSE
- Microsoft Windows 2008 Enterprise/Standard Editions; 2008 Server Core; 2003 Enterprise/Standard Editions
- Microsoft Windows XP Professional Workstation
- HP OpenVMS 8.3-1H1 (IA64) (EVA iSCSI Connectivity Option (mpx100) only)
- Sun Solaris (SPARC and x86) (EVA4000/4100/4400/4400 embedded switch/6000/6100/6400/8000/8100/8400 only)
- VMware ESX with the following guest operating systems: Windows 2003, Red Hat, and SUSE

---

**NOTE:** For the latest information on version support, see the product release notes or SPOCK at <http://www.hp.com/storage/spock>. You must sign up for an HP Passport to enable access.

---

The EVA and EVA4400 iSCSI options are compatible with all HP-supported GbE NICs for HP OpenVMS (mpx100 only), Microsoft Windows, Linux, and VMware and all standard GbE NICs supported by Apple and Sun.

### NIC Teaming

NIC Teaming is supported in single-path or multipath configurations (team failover only).

### iSCSI initiators

The EVA iSCSI Connectivity Option supports the ATTO Apple Mac OS X iSCSI Initiator, HP OpenVMS native iSCSI Initiator, Microsoft Windows iSCSI Initiator, Sun Initiator, VMware Initiator, and for Red Hat and SUSE, the bundled iSCSI driver. Contact an HP storage representative for the latest information on iSCSI initiator version support.

### iSCSI boot

iSCSI boot is supported for the following operating systems and network interface cards:

- Microsoft Windows Server 2003 SP2
- Red Hat Enterprise Linux 4 AS, Update 7, Update 6, Update 5, and Update 4
- SUSE Linux Enterprise Server 9 SP4, SP3
- SUSE Linux Enterprise Server 10, 10 SP2, SP1

QLogic QLA4052C, QLE4062C, QMH4062C iSCSI HBA is supported on the following operating systems:

- Microsoft Windows Server 2008, 2003 SP2
- Red Hat Enterprise Linux 5, Update 2, Update 1
- Red Hat Enterprise Linux 4 AS, Update 7, Update 6, Update 5
- SUSE Linux Enterprise Server 10, 10 SP2, SP1
- SUSE Linux Enterprise Server 9 SP4, SP3

## EVA and EVA4400 iSCSI Connectivity Options supported maximums

Table 198 (page 340) lists the maximums supported by the EVA and EVA4400 iSCSI Connectivity Options using the mpx100/100b bridge.

**Table 198 Supported EVA iSCSI Connectivity Option maximums**

Description	Maximum per EVA or EVA4400 iSCSI Connectivity Option
<b>Hardware</b>	
EVA storage system	1
mpx100/100b	2
<b>Configuration</b>	
Total number of iSCSI initiators	mpx100— 150 (single-path or multipath) mpx100b (EVA4400 only)— 16, 48 (license upgrade 1), 150 (license upgrade 2) (single-path or multipath) <b>Note:</b> The mpx100/100b can serve both single-path and multipath LUNs concurrently.
Total number of iSCSI LUNs	150
Total number of iSCSI targets per initiator	8

## General rules for the EVA and EVA4400 iSCSI Connectivity Options

**NOTE:** The EVA iSCSI Connectivity Option (mpx100) is supported with EVA4000/4100/4400/4400 (embedded switch)/6000/6100/6400/8000/8100/8400 and EVA3000/5000 storage systems. The EVA4400 iSCSI Connectivity Option (mpx100b) is supported with EVA4400 and EVA4400 (embedded switch) storage systems.

Use the following general rules when implementing the EVA or EVA4400 iSCSI Connectivity Option:

- Each EVA storage system can have a maximum of two mpx100 or two mpx100b bridges.
- Each EVA controller host port can connect to a maximum of two mpx100/100b Fibre Channel ports.
- Both mpx100/100b Fibre Channel ports can only connect to one EVA storage system.
- Each mpx100/100b Fibre Channel port can only connect to one EVA port.
- Each iSCSI initiator can have a maximum of eight mpx100/100b iSCSI targets.

Table 199 (page 341) describes the maximum number of servers, initiators, and LUNs supported for combined iSCSI and Fibre Channel connectivity to an EVA configured with an mpx100 or 100b.

**Table 199 Supported EVA/mpx100/100b maximums**

	EVA with Fibre Channel only	EVA with Fibre Channel and 1GbE iSCSI (mpx100)	EVA4400 with Fibre Channel and 1GbE iSCSI (mpx100b)
Maximum number of servers	256	405 with 1 EVA	271, 303, 405 with 1 EVA <sup>1</sup>
Maximum number of initiators	1,024	1,170 with 1 EVA	1,036, 1,068, 1,170 with 1 EVA <sup>1</sup>
Maximum number of LUNs <sup>2</sup>	1,023 EVA4x00/6000/6100/8000/8100 2047 EVA6400/8400	150 iSCSI 1,023 EVA4x00/6000/6100/8000/8100 2,047 EVA6400/8400	150 iSCSI 1,023 EVA4x00/6000/6100/8000/8100 2,047 EVA6400/8400

<sup>1</sup> The mpx100b supports 16 (base), 48 (license upgrade 1), and 150 (license upgrade 2) iSCSI initiators.

<sup>2</sup> For more information, see “Configuration parameters” (page 216).

## B-series iSCSI Director Blade

The B-series iSCSI Director Blade (FC4-16IP) is a gateway device between Fibre Channel targets and iSCSI initiators. This allows iSCSI initiators in an IP SAN to access Fibre Channel storage in a Fibre Channel SAN.

This section describes the following topics:

- “Blade overview” (page 341)
- “Hardware support” (page 342)
- “Software support” (page 343)
- “Scalability rules” (page 343)

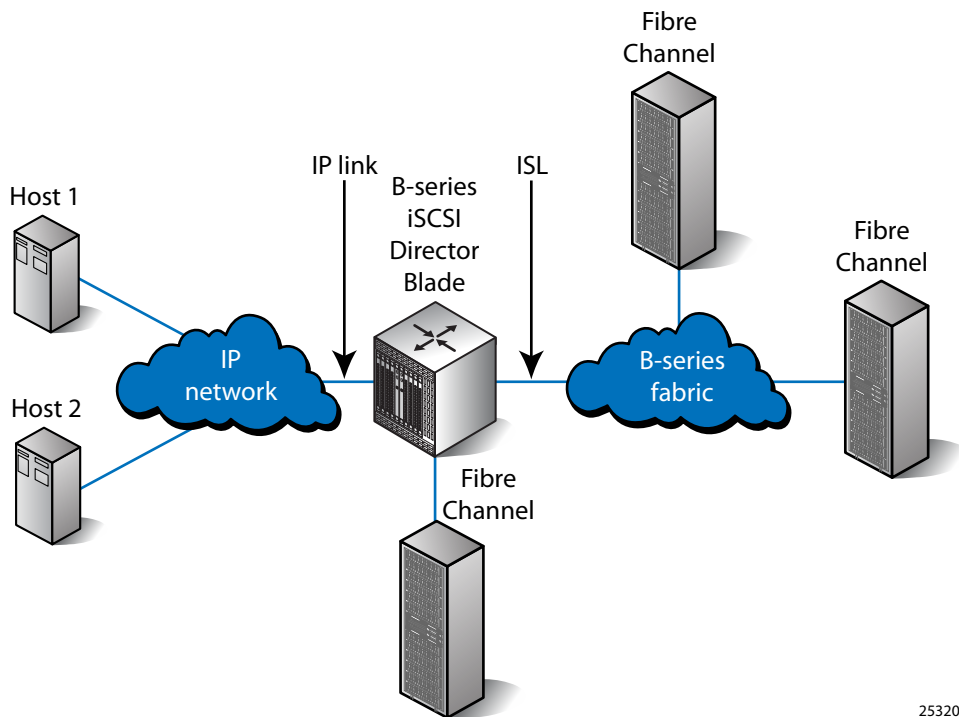
For the latest B-series documentation, see <http://h18006.www1.hp.com/storage/saninfrastructure.html>.

### Blade overview

The B-series iSCSI Director Blade provides IP hosts access to Fibre Channel storage devices. The IP host sends SCSI commands encapsulated in iSCSI PDUs to a blade port over TCP/IP. The PDUs are routed from the IP network to the Fibre Channel network and then forwarded to the target.

Figure 152 (page 342) shows a sample configuration in which a B-series iSCSI Director Blade bridges an IP network and Fibre Channel network.

**Figure 152 Fibre Channel and IP configuration with the B-series iSCSI Director Blade**



25320b

The incoming iSCSI initiators on an iSCSI port are mapped as a single iSCSI virtual initiator. The virtual initiator is presented to Fibre Channel targets as an N\_Port Fibre Channel initiator device with a WWN.

By default, the blade uses basic LUN mapping to map Fibre Channel targets to iSCSI virtual targets. This creates one iSCSI virtual target per Fibre Channel target and allows 1-to-1 mapping. The iSCSI initiators are then allowed to access the iSCSI virtual targets.

## Hardware support

This section describes the devices compatible with the B-series iSCSI Director Blade.

### Storage systems

The following storage systems are supported with the B-series iSCSI Director Blade. Contact an HP storage representative for specific support information.

- MSA1500
- EVA3000/5000
- EVA4000/6000/8000
- XP48/512
- XP128/1024
- XP10000/12000

---

**NOTE:** HP supports direct connection of storage systems to the Fibre Channel ports on the B-series iSCSI Director Blade, not to the IP ports.

---

### Fibre Channel switches

The B-series iSCSI Director Blade is supported on the B-series SAN Director 4/256, with a maximum of four blades per chassis. For more information, see [“B-series switches and fabric rules”](#) (page 83).

## Software support

This section describes the operating systems and software supported with the B-series iSCSI Director Blade.

### Operating systems and network interface controllers

The following operating systems are supported with the B-series iSCSI Director Blade and iSCSI. Each operating system's HP-supported NICs are also supported.

- Windows 2003 SP1/R2 Standard Edition, Enterprise Edition
- Windows 2003 64-bit
- Red Hat Linux 4 (32-bit and 64-bit)
- SUSE Linux Enterprise Server 9 (32-bit and 64-bit)

---

**NOTE:** The QLogic iSCSI HBA (QLA 4052c) is supported on Windows 2003, Red Hat, and SUSE.

---

### Network Teaming

The B-series iSCSI Director Blade supports Network Teaming for Linux.

### B-series management applications

The B-series iSCSI Director Blade management applications are as follows:

- B-series Fabric Manager
- CLI

### iSCSI initiators

The B-series iSCSI Director Blade supports the following iSCSI initiators:

- Microsoft Windows iSCSI Initiator
- Bundled Red Hat iSCSI Initiator
- Bundled SUSE iSCSI Initiator

Contact an HP storage representative for specific support information.

## Scalability rules

Table 200 (page 343) provides the B-series iSCSI Director Blade scalability rules.

**Table 200 B-series iSCSI Director Blade scalability rules**

Rule	Maximum
iSCSI sessions per port 64	64
iSCSI ports per FC4–16IP blade	8
iSCSI blades per switch	4
iSCSI sessions per FC4–16IP blade	512
iSCSI sessions per switch	1,024
TCP connections per switch	1,024
TCP connections per iSCSI session	2
iSCSI sessions per fabric	4,096
TCP connections per fabric	4,096
iSCSI targets per fabric	4,096

**Table 200 B-series iSCSI Director Blade scalability rules** *(continued)*

Rule	Maximum
CHAP entries per fabric	4,096
LUNs per iSCSI target	256
Members per discovery domain	64
Discovery domains per discovery domain set	4,096
Discovery domain sets	4

## B-series MP Router and iSCSI

The B-series MP Router offers Fibre Channel routing, FCIP SAN extension functionality, and iSCSI-to-Fibre Channel bridging capability.

### Hardware support

This section describes the hardware devices supported by the B-series MP Router for iSCSI.

#### Fibre Channel switches

The MP Router iSCSI mode supports any B-series switch in an MP Router backbone fabric.

The MP Router backbone fabric consists of at least one MP Router and includes all other routers and B-series switches connected to the MP Router through E\_Ports. iSCSI initiators connected to the MP Router's iSCSI ports can only access storage connected to the MP Router backbone fabric ("[iSCSI storage device connectivity for the MP Router](#)" (page 108)).

Edge fabrics consist of all Fibre Channel switches and their attached storage connected to the MP Router through EX\_Ports. Edge fabrics are not accessible to iSCSI initiators connected to the MP Router's iSCSI ports ("[Fibre Channel routing storage device connectivity for the MP Router](#)" (page 109)).

#### Storage systems

MP Router iSCSI devices must connect directly to the backbone fabric of the router. The supported arrays are as follows:

- MSA1000
- MSA1500
- EVA3000/5000
- EVA4000/6000/8000
- XP storage systems

### Software support

MP Router iSCSI supports single-path devices only. Multipath software is not supported.

#### Operating systems and network interface controllers

The MP Router supports the following operating systems:

- Microsoft Windows 2000 SP3, SP4
- Microsoft Windows 2003 Standard Server and Enterprise Editions

The MP Router in iSCSI mode is compatible with all HP-supported NICs for Windows.

#### Network Teaming

MP Router iSCSI supports HP Network Teaming with Windows 2000 and Windows 2003.



## iSCSI initiators

MP Router iSCSI supports Microsoft iSCSI initiator 1.06.

## Configuration rules

Table 201 (page 345) describes the iSCSI configuration rules for the MP Router.

**Table 201 MP Router iSCSI configuration rules**

Rule	Maximum
Maximum number of iSCSI ports per MP Router	12
Maximum number of TCP sessions per MP Router	96
Number of TCP sessions per port	8
Maximum number of targets per host	1

The Windows iSCSI Initiator software currently supports a maximum of one target per host when that target is accessed through the MP Router in iSCSI mode. This prevents the use of multipath software.

## C-series iSCSI

The C-series IP Storage Services Modules, 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, MDS 9216i Multilayer Fabric switch, and MDS 9222i Multiservice Fabric switch provide iSCSI capabilities. These devices allow IP hosts to access Fibre Channel storage using iSCSI.

The IP Storage Services Modules, 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, 9216i switch, and MDS 9222i Multiservice Fabric switch provide transparent iSCSI routing by default. IP hosts using iSCSI can access targets on the Fibre Channel network.

This section describes the following topics:

- “Modules overview” (page 345)
- “Hardware support” (page 346)
- “Software support” (page 347)
- “Configuration rules” (page 348)

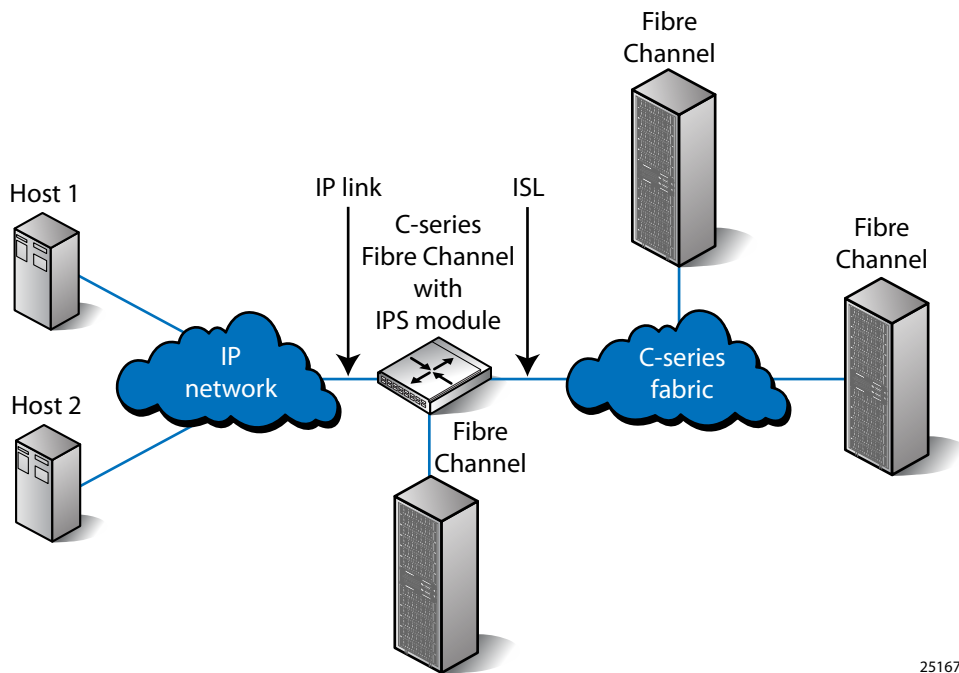
For the latest C-series documentation, see <http://h18006.www1.hp.com/storage/saninfrastructure.html>.

## Modules overview

The C-series IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, 9216i switch, and MDS 9222i Multiservice Fabric switch provide IP hosts access to Fibre Channel storage systems. The IP host sends SCSI commands encapsulated in iSCSI PDUs to a module port over a TCP/IP connection. These commands are routed from the IP network to the Fibre Channel network and then forwarded to the target.

Figure 153 (page 346) shows a sample configuration where an IP module bridges an IP network and a Fibre Channel network.

**Figure 153 C-series Fibre Channel and IP configuration with the IP module**



In addition to presenting Fibre Channel targets to iSCSI hosts, the modules also present each iSCSI host as a Fibre Channel host (in transparent mode). The iSCSI host appears as an HBA to the Fibre Channel storage device. The storage device responds to each IP host as if a Fibre Channel host were connected to the Fibre Channel network.

## Hardware support

This section describes the devices compatible with the C-series IP Storage Services Modules, 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, 9216i switch, and MDS 9222i Multiservice Fabric switch.

## Storage systems

This section lists the storage array iSCSI support for the C-series modules. Storage arrays and array options like Continuous Access do not have the same iSCSI support across all operating systems. Contact an HP storage representative for specific support information.

The following storage systems are supported with iSCSI:

- MSA1000/1500
- EVA3000/5000
- EVA4000/6000/8000
- EVA4100/6100/8100
- VA7100/7110
- VA7400/7410
- XP48/512
- XP128/1024
- XP10000/12000
- XP20000/24000

**NOTE:** HP supports direct connection of storage arrays to the Fibre Channel ports on the 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, MDS 9216i, and MDS 9222i Multiservice Fabric. HP does *not* support direct connection of storage arrays to the IP Storage Services Modules (IPS-4, IPS-8) or to the IP ports on the 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, MDS 9216i, or MDS 9222i Multiservice Fabric.

### Fibre Channel switches

The C-series IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Services Module, and 18/4 Multiservice Module support the HP C-series switches listed in “C-series switches and fabric rules” (page 115).

Table 202 (page 347) describes the C-series switches that support iSCSI.

**Table 202 C-series switches that support iSCSI**

Switch	Maximum number of Fibre Channel ports	Supported IP modules
MDS 9216	48	IPS-4, IPS-8, and 14/2
MDS 9216A	64	IPS-4, IPS-8, 14/2, and 18/4
MDS 9216i	62	Embedded 2 ports and IPS-4, IPS-8, 14/2, and 18/4
MDS 9222i Multiservice Fabric	66	IPS-8, 18/4
MDS 9506	192	IPS-4, IPS-8, 14/2, and 18/4
MDS 9509	336	IPS-4, IPS-8, 14/2, and 18/4
MDS 9513	528	IPS-4, IPS-8, 14/2, and 18/4

### Software support

This section describes the operating systems and related software compatible with the C-series IP Storage Services Modules (IPS-4, IPS-8), 14/2 Multiprotocol Service Module, 18/4 Multiservice Module, MDS 9216i switch, and MDS 9222i Multiservice Fabric switch.

#### Operating systems and network interface controllers

The following operating systems are supported with C-series modules and iSCSI. Each operating system's HP-supported NICs are also supported with iSCSI.

- Windows 2000 Server, Advanced Server
- Windows 2003 Standard Edition, Enterprise Edition
- Windows 2003 64-bit
- Windows 2003 x64
- Red Hat Enterprise Linux AS, ES, WS
- SUSE Linux Enterprise Server
- IBM AIX
- Sun Solaris
- VMware ESX

#### HP Network Teaming

Windows 2000 and Windows 2003 support HP Network Teaming.

## C-series management applications

C-series management applications are as follows:

- Cisco Fabric Manager
- Cisco Device Manager
- CLI

## iSCSI initiators

C-series modules support these iSCSI software initiators:

- Microsoft Windows iSCSI Initiator
- Red Hat iSCSI bundled Initiator
- Red Hat SourceForge iSCSI Initiator
- SUSE iSCSI bundled Initiator
- SUSE SourceForge iSCSI Initiator
- IBM AIX native iSCSI Initiator
- Sun Solaris native iSCSI Initiator
- VMware native iSCSI Initiator

## Configuration rules

This section describes the iSCSI limits and rules when using the IP Storage Services Modules, 14/2 Multiprotocol Services Modules, and 18/4 Multiservice Modules in the following:

- MDS 9506, MDS 9509, and MDS 9513 Director switches
- MDS 9216, MDS 9216A, MDS 9216i, and MDS 9222i Fabric switches
- Embedded 14/2 Multiprotocol Services Module in the MDS 9216i Fabric switch
- Embedded 18/4 Multiservice Module in the MDS 9222i Multiservice Fabric switch

Multipathing is supported only in a Windows environment using the Microsoft MPIO iSCSI driver. HP Secure Path is not supported with iSCSI initiators. Contact an HP storage representative for specific iSCSI support information.

Without multipathing capabilities, the iSCSI initiator can only access one path of the storage controller, which disables controller failover protection.

The IP Storage Services Modules, 14/2 Multiprotocol Services Module, 18/4 Multiservice Module, 9216i switch, and MDS 9222i Multiservice Fabric switch are supported on fabrics constructed with C-series switches. See “[C-series switches and fabric rules](#)” (page 115) for the latest C-series fabric rules.

[Table 203](#) (page 348) describes C-series module iSCSI limits.

**Table 203 C-series iSCSI limits**

C-series iSCSI limits	Maximum
Number of initiator/target pairs per port	500
Number of active LUNs per initiator/target pair	256
Number of initiator/target pair/LUN combinations per GbE port	1,200

The following examples show maximum configurations for initiator/target pairs:

- 500 iSCSI initiators, each connecting to one target (storage controller port)
- 100 iSCSI initiators, each connecting to five targets
- 50 iSCSI initiators, each connecting to eight targets and 100 iSCSI initiators, each connecting to one target

The maximum of 500 TCP connections (initiator/target pairs) with 256 LUNs per connection yields 128,000 possible LUNs. Simultaneous access to 128,000 LUNs via a single 100 Mb Ethernet port may provide reduced performance.

The total bandwidth required for an Ethernet port for iSCSI storage should be based on the sum of the individual requirements for each connection.

## HP ProLiant Storage Server iSCSI Feature Pack

This section describes the following topics:

- “Overview” (page 349)
- “HP ProLiant Storage Server iSCSI Feature Pack support” (page 349)
- “HP ProLiant Storage Server iSCSI license upgrade options” (page 350)
- “Designing a Microsoft Exchange solution with iSCSI Feature Pack” (page 351)
- “Sample iSCSI NAS Microsoft Exchange Server 2003 configuration” (page 353)

### Overview

HP ProLiant Storage Servers can be configured as iSCSI targets using the HP ProLiant Storage Server iSCSI Feature Pack.

The HP iSCSI storage server solution facilitates centralized management, backup, and scalability by integrating services for:

- Files
- Printing
- Email
- Databases

Existing Ethernet infrastructure provides low-cost storage consolidation.

The HP iSCSI Feature Pack (T3669A only) includes the HP ProLiant Application Storage Manager, which provides storage management for HP NAS servers hosting Microsoft Exchange 2003 storage groups. It reduces the time and training required to set up and monitor email stores.

The HP iSCSI storage server is ideal for businesses that want to consolidate storage and that use applications such as Microsoft Exchange.

For more information, see <http://h18006.www1.hp.com/storage/storageservers.html>.

### HP ProLiant Storage Server iSCSI Feature Pack support

This section describes support for the HP ProLiant Storage Server iSCSI Feature Pack.

#### Hardware support

The following storage servers support the HP iSCSI Feature Pack:

- HP ProLiant DL100 Storage Server
- HP ProLiant ML110 Storage Server
- HP ProLiant ML350 G4 Storage Server
- HP ProLiant ML370 G4 Storage Server

- HP ProLiant DL380 G4 Storage Server (Base, External SCSI, and External SATA models, SAN Storage model Gateway Edition only)
- HP StorageWorks NAS 500s
- HP StorageWorks NAS 1200s
- HP StorageWorks NAS 1500s
- HP StorageWorks NAS 2000s
- HP StorageWorks NAS 4000s (Gateway Edition only)

### Application support

The following host applications support the HP iSCSI Feature Pack:

- Microsoft Exchange Server 2000
- Microsoft Exchange Server 2003
- Microsoft SQL Server 2000
- Microsoft SQL Server 2003
- Oracle Database 9i and 10g

### Management software support

You can perform management functions with the HP Storage Server GUI.

### iSCSI Initiator support rules

HP iSCSI Feature Pack support rules follow:

- The Microsoft iSCSI Initiator (32-bit version) is supported.
- Up to 50 simultaneous initiators are supported.
- Initiators running Microsoft XP Home Edition or Microsoft XP Professional Edition have not been qualified.
- Hardware initiators may work but are not currently supported.

---

**NOTE:** Linux and HP-UX initiators have not been fully qualified with this solution. Hardware initiators are not currently supported.

---

## HP ProLiant Storage Server iSCSI license upgrade options

The HP ProLiant Storage Server iSCSI Feature Pack has three licensed options:

- “Snapshot” (page 350)
- “Clustering” (page 351)
- “Direct Backup” (page 351)

### Snapshot

The Snapshot option is an upgrade license for HP ProLiant Storage Server iSCSI Feature Pack. The Snapshot option:

- Works with Microsoft iSCSI initiators only.
- Prevents accidental deletions, file corruption, and virus attacks.
- Pauses application hosts running Microsoft Exchange, SQL Server, or Oracle Database to ensure data integrity.

- Allows delta snapshots using Microsoft VSS interface, and performs automatic delta snapshots of application hosts to reduce potential data loss.
- Offers several application-specific licensed agent options:
  - Microsoft Visual SourceSafe
  - Microsoft Exchange
  - Microsoft SQL
  - Oracle Database (for a single Microsoft iSCSI initiator)

### Clustering

Clustering is an upgrade license for the HP ProLiant Storage Server iSCSI Feature Pack (Gateway Edition only). The clustering option:

- Activates two-node iSCSI target capability using MSCS
- Eliminates a single point of failure with a dual network connection to the IP network and a dual I/O channel to each storage device

### Direct Backup

The Direct Backup option is an upgrade license for HP ProLiant Storage Server iSCSI Feature Pack. The Direct Backup option:

- Works with Microsoft iSCSI initiators only
- Facilitates centralized, zero-impact backup
- Allows administrators to use their preferred backup software for centralized backup and recovery of application data directly from the iSCSI storage server

## Designing a Microsoft Exchange solution with iSCSI Feature Pack

A Microsoft Exchange solution requires that you configure the network, host, and storage systems for iSCSI NAS. The HP ProLiant Storage Server iSCSI Feature Pack provides iSCSI functionality on a Windows Storage Server (NAS device). Exchange Server 2003 also requires an iSCSI initiator.

### Network design

Existing IP networks may not be suitable for iSCSI storage support. Evaluate traffic on these networks to determine if there is adequate capacity to meet storage requirements.

HP recommends a dedicated GbE network for accessing the Windows iSCSI NAS Storage Server. This provides adequate performance and data security. You can also use IPsec to secure the connection on a public, unsecured network, with decreased performance.

The distance between the Exchange server and the iSCSI NAS Storage Server may affect performance. You must check the maximum supported distances of the network devices. The maximum distance varies based on the cable type and specifications.

### Hardware selection

The Windows Server Catalog is available at <http://www.windowsservercatalog.com/>.

It lists iSCSI hardware components that are qualified under the Designed for Windows Logo program, Exchange Server 2003 and Exchange Server 2000.

## Exchange storage design

Important criteria for Exchange storage design include:

- Isolation of Exchange transaction logs from databases
- Selection of optimum RAID level for performance and fault tolerance
- Write-back caching for hardware RAID controller performance

### Separate volumes for logs and databases

HP recommends separate volumes for Exchange transaction logs and databases, to ensure data protection and efficiency. Transaction log access is mostly sequential writes does database access is random read/write.

The Exchange server internal storage can hold the Exchange transaction logs while the HP Storage Server with the HP iSCSI Feature Pack holds the Exchange databases.

The Exchange server internal storage and the iSCSI NAS Storage Server have comparable transaction log performance. If the transaction logs are stored on the Exchange server, you must recover them manually if the server fails. Regaining access to the log drives requires installing the logs on a backup Exchange server.

---

**NOTE:** Store transaction logs on a RAID 1 mirror pair array volume, or, for additional disk space, on four or more disks on a RAID 1+0 (striped mirror).

---

### RAID level

For database volumes, the choice of RAID protection on the disk arrays is a trade-off between maximum storage and performance. For the same number of disk spindles, RAID 5 provides data protection and maximum storage, and RAID 1+0 provides the best performance but with less storage.

Six spindles in a RAID 1+0 array provide greater performance than six spindles in a RAID 5 array for a given number of Exchange mailboxes. RAID 1+0 is preferred for the database volume. Using six 36 GB drives in a RAID 1+0 array provides ample storage and performance for 1,000 100 MB mailboxes.

Designing storage arrays with the large disk drives (146 GB or larger) requires caution. Although a few large disks can provide the required amount of database storage, the reduced spindle count will decrease I/O performance.

### Write-back caching

Write-back caching increases performance for transaction logs on the server array and for databases on the storage array. On array controllers with battery-backed write cache (such as the Smart Array 5i Plus and later) the write-cache percentage should be set to 100%.

### Recommendations

HP recommends that you:

- Place the Exchange log files and database files on separate RAID 1+0 RAID sets.
- Place the Exchange log files on Exchange server disks.
- Use hardware RAID controllers with write-back caching.

### Supported load with Exchange

The performance required by the average email user determines the storage design. The average load is multiplied by the number of users to find the storage requirement. Conversely, the capabilities of an existing system can determine the maximum number of users.

To calculate the average I/O per user in an Exchange environment, the PERFMON object's disk-transfers-per-second value is divided by the number of active connections. The storage capacity



calculated from the average I/O needs an additional safety factor to maintain performance during peak periods.

In practice, the maximum number of users is less than the calculated value when:

- Users increase the size of their mailboxes.
- Services such as antivirus scanners or content indexers are added to the Exchange server.

A medium-sized user profile provides a 60 MB mailbox, and a large profile provides a 100 MB mailbox. Larger mailboxes affect both storage sizing and performance, and are disproportionately more difficult for Exchange to manage.

## Sample iSCSI NAS Microsoft Exchange Server 2003 configuration

This section summarizes results from an HP-tested iSCSI NAS configuration. Details of this test are available at [http://storage.corp.hp.com/Document\\_Storage/whitepapers/new\\_library/5982-9399EN\\_Exchange\\_2003\\_HP\\_iSCSI\\_and\\_WSS\\_NAS\\_WP.pdf](http://storage.corp.hp.com/Document_Storage/whitepapers/new_library/5982-9399EN_Exchange_2003_HP_iSCSI_and_WSS_NAS_WP.pdf).

### HP ProLiant DL380 G4 Storage Server configuration

With up to six internal drives and fault-tolerant options for fans and power supplies, an HP ProLiant DL380 server acts as a high-capacity Exchange server. Other HP ProLiant servers can support 900 to 2,000 Exchange mailboxes.

When configuring the server and storage, ensure that you have:

- Sufficient RAM (1 GB) on the server.
- Enough disks for the Exchange database volume, especially if you use RAID 5 RAIDsets or ADG. A larger number of disks improves database volume performance.

The Exchange server internal RAID storage holds the transaction logs. HP recommends storing the transaction logs on a RAID array accelerated with a write-back cache, with battery backup. The transaction logs can be stored on an iSCSI NAS RAID array.

For the sample configuration, the server includes a Smart Array 5i RAID controller option kit since the write cache for this controller can be set at 100%, increasing performance.

RAID 1+0 is preferred for the database volume on the storage server. In the sample iSCSI NAS storage configuration:

- The DL380 G4 Storage Server uses a Smart Array 6402 and Smart Array 6404 controller to access the disks for its RAID array.
- 56 disks are available for RAID configurations. These disks sit in four rack-mountable HP StorageWorks Modular Smart Array 30 (MSA30) disk enclosures with 14 drives each.
- A 40-disk RAID 1+0 configuration consists of an array of 10 disks in each of the four enclosures.

The Exchange server accesses its database on the iSCSI NAS RAID 1+0 volume.

---

# 19 Network Attached Storage

This chapter describes the HP ProLiant Storage Server family of NAS products. It describes the following topics:

- “Storage server and SAN integration” (page 354)
- “HP ProLiant Storage Servers” (page 354)

---

**NOTE:** The HP ProLiant Storage Server family was formerly called the HP StorageWorks NAS family. Legacy NAS products use the HP StorageWorks NAS brand.

For information about iSCSI NAS configurations, see “HP ProLiant Storage Server iSCSI Feature Pack support” (page 349).

---

## Storage server and SAN integration

Storage servers combine NAS and SAN in a common, networked storage pool that allows you to choose file-level (NAS) or block-level (SAN) access. Storage server integration in a heterogeneous SAN environment provides:

- Support for multiple operating systems
  - Data centralization and consolidation
  - Cost reduction through improved disk utilization
  - Increased storage capacity and efficiency
- Integration of storage networks with existing hardware
  - Improved manageability
  - Less interruption to production systems
  - Reliable application integration
  - Reduction in total cost of ownership
- Enterprise-wide file sharing
  - Reduction in file duplication
  - Increased storage capacity
- Reliable and efficient access to centrally managed data
  - Uninterrupted business operations
  - Single application server interface

## Configuration and zoning rules

Storage servers follow the same configuration and zoning rules as the Windows 2003 operating system. To determine the configuration and zoning rules for a specific storage server, see “Heterogeneous server rules” (page 159).

## HP ProLiant Storage Servers

The HP ProLiant Storage Servers provide scalability, flexibility, and cost-effective management of storage resources. HP offers the following storage servers:

- HP ProLiant DL380 G4
- HP ProLiant DL580 G2
- HP ProLiant DL585

---

**NOTE:** The DL380 G4 was formerly called NAS 4000s; the DL580 G2 was called NAS 9000s.

---

## Features

The HP ProLiant Storage Servers offer the following features:

- Multi-protocol file serving (Windows, UNIX, Linux, FTP, HTTP, AppleTalk)
- Native integration into Microsoft Active Directory
- Storage quotas (user and directory)
- Snapshots
- High-availability clustering through MSCS
- Simplified web-based management
- Systems Insight Manager
- Redundant hardware components
- iLO connectivity
- 10/100/1000 Ethernet (TOE is optional.)
- Standard 3-3-3 warranty (HP Care Pack is optional.)

## Software

The DL380 G4, DL580 G2, and DL585 Storage Servers have Windows Storage Server 2003 Enterprise Edition preinstalled and preconfigured.

HP offers the following optional software packages:

- HP ProLiant Storage Server iSCSI Feature Pack
- Data replication (HP OpenView Storage Mirroring)
- Backup agents
- Antivirus software

The HP iSCSI Feature Pack provides iSCSI target functionality to HP storage servers. It enables hosting of file, print, and application storage services on a single platform.

For more information, see “HP ProLiant Storage Server iSCSI Feature Pack support” (page 349) or the HP ProLiant Storage Server iSCSI Feature Pack website <http://h18006.www1.hp.com/products/storage/software/inas/index.html>.

## Hardware

The DL380 G4, DL580 G2, and DL585 are supported in a heterogeneous open SAN as both standalone and clustered servers. They follow the Windows 2003 SAN hardware rules.

For additional information, see the following QuickSpecs:

- DL380 G4: <http://h18006.www1.hp.com/products/storageworks/dl380storageserver/index.html>
- DL580 G2: <http://h18006.www1.hp.com/products/storageworks/dl580storageserver/specifications.html>
- DL585: [http://h18006.www1.hp.com/products/quickspecs/11902\\_na/11902\\_na.html](http://h18006.www1.hp.com/products/quickspecs/11902_na/11902_na.html)

## Fabric rules

Storage servers are supported in SAN fabrics with B-series, C-series, or M-series switches. For more information about fabric rules, see [“B-series switches and fabric rules” \(page 83\)](#), [“C-series switches and fabric rules” \(page 115\)](#), or [“M-series switches and fabric rules” \(page 140\)](#).

## Storage rules

The DL380 G4, DL580 G2, and DL585 support Fibre Channel storage connectivity using the following arrays:

- XP
- VA
- EVA
- MSA

Supported storage system firmware versions are the same as those listed in this guide for Windows 2003. For more information, see the storage system documentation or contact your HP representative.

---

# Part V Storage security, best practices, and support information

The following chapters describe storage security and SAN best practices:

- “Storage security” (page 358)
- “Best practices” (page 372)
- “Support and other resources” (page 393)

## 20 Storage security

This chapter describes storage security best practices. It describes the following topics:

- “Storage security threats” (page 358)
- “Storage security compliance” (page 359)
- “Security technologies” (page 359)
- “HP security strategy” (page 361)
- “Storage security best practices” (page 363)
- “Assessing security risks” (page 363)
- “HP storage security solutions” (page 364)

### Storage security threats

Securing SAN environments has become an increasingly important aspect of data security. IT organizations face many security threats and must comply with numerous industry and government regulations. In the past, IT organizations accepted that authentication issues were handled by the network architecture; they were not responsible for SAN security.

The NSA IATF defines five security attack classes that you should consider when defining your solution (Table 204 (page 358)).

**Table 204 Security attack classes**

Attack class	Description
Passive	Attacks that can disclose information to an attacker. Passive attacks include: <ul style="list-style-type: none"><li>• Analyzing traffic</li><li>• Monitoring unprotected communications</li><li>• Decrypting weakly encrypted traffic</li><li>• Capturing authentication information (passwords)</li></ul> An example of a passive attack is the disclosure of information such as credit card numbers and passwords.
Active	Attacks that can disclose information, deny service, or modify data. Active attacks include: <ul style="list-style-type: none"><li>• Attempting to circumvent or break protection features</li><li>• Introducing malicious code</li><li>• Stealing or modifying information</li><li>• Attacking a network backbone</li><li>• Exploiting in-transit information</li><li>• Penetrating an enclave</li><li>• Attacking when a remote user attempts to connect to an enclave</li></ul>
Close-in	Attacks by an unauthorized user who is in close physical proximity to networks, systems, or facilities. The user may attempt to gather or modify information, or deny authorized users access to information.
Insider	Unauthorized attacks by an authorized user can be malicious or nonmalicious. Malicious attackers can: <ul style="list-style-type: none"><li>• Eavesdrop</li><li>• Steal or damage data</li></ul>

**Table 204 Security attack classes** *(continued)*

Attack class	Description
	<ul style="list-style-type: none"> <li>• Use data for fraudulent purposes</li> <li>• Deny authorized users access</li> </ul> <p>Nonmalicious attacks can result from:</p> <ul style="list-style-type: none"> <li>• Carelessness</li> <li>• Lack of knowledge</li> <li>• Circumventing security for nonmalicious purposes to perform tasks</li> </ul>
Distribution	Attacks due to modifications to hardware or software made at the factory or during distribution. Distribution attacks can insert malicious code in a product, which can allow future unauthorized access to the system.

## Storage security compliance

Compliance ensures that a storage system meets specific criteria established by law or regulation. Retention of electronic records is mandated by statutory and regulatory law.

Data security regulations are enacted by international governments and U.S. federal and state governments. All storage systems must comply with local regulations. [Table 205 \(page 359\)](#) lists some of the U.S. and international security regulations.

**Table 205 U.S. and international security regulations**

U.S. federal and state regulations	International regulations
<ul style="list-style-type: none"> <li>• Sarbanes-Oxley (SOX) Act of 2002</li> <li>• Gramm-Leach-Bliley Act (GLBA) of 1999</li> <li>• Securities and Exchange Commission Act (SEC) rules 17a-3 and 17a-4</li> <li>• Department of Energy (DOE) 10 CFR 600.153 Retention and access requirements for records</li> <li>• California Data Security Act (SB 1386/AB 1950)</li> <li>• New York Regulation 173 Standards for safeguarding customer information</li> </ul>	<ul style="list-style-type: none"> <li>• European Union Data Protection Directive of 1995</li> <li>• Canada: Personal Information Protection and Electronic Documents Act (PIPEDA)</li> <li>• Australia: Privacy Act 1988</li> <li>• Japan: Personal Information Protection Act</li> <li>• UK: Data Protection Act 1998</li> <li>• New Zealand: Privacy Act 1993</li> </ul>

## Security technologies

This section describes security technologies for IP SAN, Fibre Channel SAN, and encryption.

### IP SAN security technologies

IP SAN technologies includes NAS, iSCSI, and FCIP. IP SAN security is achieved through the following:

- CHAP
- IPsec

#### CHAP

CHAP uses a three-way handshake to ensure validity of remote clients. It is more secure than the PAP. A summary of the CHAP process follows:

1. When the server is first connected, it sends a challenge message to the peer.
2. The peer responds by sending a value generated by a one-way hash function.
3. The server compares this value to its own generated value.

4. If the values match, the connection is allowed to continue; if they do not match, the connection is terminated.
5. To ensure the validity of the peer, the server sends challenge messages at random intervals and changes the CHAP identifiers frequently.

## IPsec

IPsec uses an open-standards framework to protect data transmission over IP networks. It uses cryptographic security services.

IPsec supports:

- Network-level peer authentication
- Data-origin authentication
- Data integrity
- Data encryption
- Replay protection

Microsoft bases its IPsec implementation on the standards developed by the IETF IPsec working group.

## Fibre Channel SAN security technologies

Fibre Channel SAN security is achieved through the FC-SP.

### FC-SP

FC-SP protects in-transit data—It does not protect data stored on the Fibre Channel network. FC-SP is a project of the Technical Committee T11, within the International Committee for Information Technology Standards, which is responsible for developing Fibre Channel interfaces (see <http://www.t11.org>).

FC-SP uses:

- Authentication of Fibre Channel devices (device-to-device authentication)
- Cryptographically secure key exchange
- Cryptographically secure communication between Fibre Channel devices

## Encryption security technologies

Encryption security is achieved through the DES, AES, and key management.

### Data Encryption Standard

DES is a block cipher designed for use in symmetric cryptography, which encrypts data in 64-bit blocks and uses a key length of 56 bits. It uses a 64-bit key, but every eighth bit is ignored. These extra bits can be used for other purposes, such as a parity check to ensure that the key is error free.

The DES cipher consists of the following process:

1. Performing an initial permutation
2. Breaking the block into right and left halves (32 bits each, followed by 16 key-dependent rounds on each half)
3. Rejoining of the halves
4. Performing the final permutation (reverse of the initial permutation)

Two common DES cipher modes are as follows:



- **ECB**—Each block of the message is encrypted independently.
- **CBC**—Each plaintext block uses an Exclusive-OR operation with the previous cipher text block before encryption.

## Advanced Encryption Standard

AES is a block cipher designed for use in symmetric cryptography, which encrypts data in 128-bit blocks. AES can use a key size of 128, 192, or 256 bits. The number of rounds varies by the key length (for example, 10, 12, or 14 rounds for key sizes 128, 192, or 256, respectively). The processing in each round is more efficient than DES and is better suited to high-speed parallel operations. A subkey step using an XOR operation, followed by a MixColumns step, occurs before the rounds are performed.

AES has equivalent modes to the ECB and CBC modes for DES. AES also has a counter mode in which a sequence number uses an Exclusive-XOR operation with the plaintext before processing; the sequence number is incremented for use with the next block.

## Key management

Successful key management is the most important yet most difficult aspect of a cryptographic system because it often requires coordination between departments and users, and the establishment and enforcement of strict system policies. You must ensure the generation, storage, exchange, verification, replacement, and destruction of keys.

## Organizational security policies

Organizational security policies are high-level statements that define the data protection requirements, which are driven by business needs. Auditing and reporting policies are added to the security policies, and the business policies are then mapped to the security policies.

## HP security strategy

This section describes the HP Secure Advantage, the HP security strategy.

## HP Secure Advantage

HP Secure Advantage allows you to combine HP security products. The Secure Advantage portfolio ensures secure automation, optimization, and acceleration of your infrastructure with proper validation to reduce risk and improve business outcomes. HP provides solutions in information security, identity management, key management, and compliance to ensure your enterprise security. Secure Advantage builds on these existing security technologies to create manageable methods for you to leverage encryption and key management. This allows you to protect your resources and validate compliance with government and industry regulations.

Security is an important aspect of the HP Adaptive Infrastructure, which provides the platform for the next-generation data center. Secure Advantage integrates with Adaptive Infrastructure enablers, such as IT systems and services, power and cooling, virtualization, and automation.

The Secure Advantage portfolio considers three aspects to ensure storage security:

- Resource protection
- Data protection
- Security validation

## Resource protection

Resource protection is important to your security strategy. Using trusted platforms, you improve availability and provide protection for networks, software, and database management systems. Access control in a trusted and hardened infrastructure minimizes disruptions due to security breaches.

## Access control

Access control prevents unauthorized use of network resources and unauthorized disclosure or modification of data (for example, preventing users from logging in to local workstations or limiting the use of dial-in modems). Access control is a set of controls: confidentiality, integrity, availability, and limiting access to network resources. These depend on the successful prevention of unauthorized access to services or information.

Important elements of access control include:

- **Identification**—Identifies an entity (user, process, or role associated with multiple users)
- **Authorization**—Determines the access rights of an entity (with a level of assurance)
- **Authentication**—Authenticates a user or process
- **Enforcement**—Applies access-control decisions, which provides protection

## Data protection

Data protection is important for all data states: at-rest, in-transit, and in-use. Use encryption and identity management in conjunction with other proactive techniques, such as security event management and information management.

Data protection consists of the following:

- Confidentiality
- Data integrity
- Data availability
- Nonrepudiation

### Confidentiality

Confidentiality prevents disclosure of all data, regardless of its state (at-rest, in-transit, or in-use). Confidentiality needs vary depending on the amount and type of data, transit and storage locations, and sensitivity of the end-user identity.

Important elements of confidentiality include:

- **Data encryption**—Invokes mechanisms that act in response to characteristics of the data, not in response to a threat.
- **Data separation**—Provides separate paths for data or processing. The level of security for data separation depends on the trust level associated with the system. Data separation ensures confidentiality by preventing data from reaching unauthorized users.
- **Traffic separation**—Adds meaningless random information and hides network-layer addresses. Traffic separation ensures confidentiality by making it difficult to determine data characteristics, such as frequency and traffic-flow destinations.

### Data integrity

Data integrity prevents unauthorized modification or destruction of data and ensures nonrepudiation and authenticity. Recording all changes to data enables the detection and notification of unauthorized modifications.

Data integrity has two types of data:

- **Single-unit data**—Applied to a single piece of data
- **Data stream**—Applied to all PDUs

### Data availability

Data availability ensures reliable access to data and information services for authorized users in the SAN. You must protect your data from attacks, unauthorized use, and routine failures.

## Nonrepudiation

Nonrepudiation ensures that all parties in a transaction are authenticated and verifies that they participated in the transaction. Storage technologies are tied closely with data and are often the last line of defense against attacks.

## Security validation

Security validation establishes a secure audit trail across your organization. The audit trail serves as proof of compliance for internal and external audits with real-time alerts. Validation is accomplished using encryption, key management, and identity management, which creates an integrated compliance solution across the organization.

To ensure compliance, every process you use must be repeatable, have demonstrated control points (with documented responsible personnel), and include a tamper-proof audit tracking system.

## Storage security best practices

To simplify storage security, the SNIA SSIF has developed the following security elements:

- **Storage system security**—Secures embedded operating systems and applications. Integrates with IT and security infrastructure, such as external authentication services, centralized logging, and firewalls.
- **SRM**—Securely provisions, monitors, tunes, reallocates, and controls storage resources to ensure storage and retrieval of data.
- **Data in-flight**—Protects the confidentiality, integrity, and availability of data as it is transferred across the SAN, LAN, or WAN. This may also include traffic management.
- **Data at-rest**—Protects the confidentiality, integrity, and availability of data stored on servers, storage arrays, NAS appliances, tape libraries, and other media. The measures required depend on the type of risk you are managing.
- **Compliance validation**—Proof of compliance is required by government and industry regulations. You must establish control points that ensure repeatable processes, assignment of responsibilities, and role separation. You must be able to prove that policies are being enforced for internal and external audits.

## Assessing security risks

This section describes best practices for assessing and addressing security risks.

## Managing organizational risks

Managing organizational risks involves the following actions:

- Protecting IT resources
- Protecting data in all states (at-rest, in-transit, or in-use)
- Providing validation to internal and external auditors

The HP Secure Advantage solution addresses these security issues using a suite of integrated products. Integration of encryption and key management technologies with identity management in a hardened infrastructure ensures that the correct data is delivered to the intended users. Secure Advantage provides the best layered end-to-end security approach with identity management at the network, system, service, and application layers. It ensures a robust and proactive security framework.

## Data security implementations

Data security implementations are categorized as follows:

- **Storage network**—Consists of switches, appliances, and cables. Switches and appliances come with support to protect themselves. The storage network components support key management, encryption services, and authentication of server and storage arrays.
- **Servers**—Consists of hardware, operating systems, interface cards (NICs and HBAs), and applications (also known as hosts). Each component comes with support for protecting itself. The interfaces cards support authentication and secure tunnel.
- **Storage arrays**—Consists of groups of disks or tapes that use a management application, which protects the resources through authentication. Storage arrays will support native encryption in the future.

## HP storage security solutions

This section describes HP storage security solutions for the following products:

- [“C-series Storage Media Encryption” \(page 364\)](#)
- [“C-series SAN-OS security” \(page 365\)](#)
- [“C-series IP SAN security” \(page 366\)](#)
- [“B-series Encryption Switch and Encryption FC Blade security” \(page 366\)](#)
- [“B-series Fabric OS security” \(page 367\)](#)
- [“Key management” \(page 370\)](#)

## C-series Storage Media Encryption

SME is a standards-based encryption solution for heterogeneous and virtual tape libraries. SME is managed with the Cisco Fabric Manager web client and a command-line interface, which supports unified SAN management and security provisioning. SME is a comprehensive network-integrated encryption service with key management that works transparently with new and existing SANs. This solution has advantages over competitive solutions, such as:

- Supports nondisruptive installation and provisioning. You do not need to rewire or reconfigure your SAN.
- Encryption engines are integrated on the MDS 9000 18/4-port Multiservice Module (MSM-18/4) and the MDS 9222i Multiservice Fabric Switch. You do not need to purchase and manage additional switch ports, cables, and applications.
- All VSAN traffic can be encrypted. This enables automated load balancing through network traffic management across multiple SANs.
- No additional software is required for key and user management or provisioning. SME is integrated with the Cisco Fabric Manager, which reduces operating expenses.

### Features

Management features of the Cisco Fabric Manager are as follows:

- Transparent fabric service
- Encryption
- Security roles
- Key management
- Clustering
- Fibre Channel redirect
- Host-based discovery for provisioning tapes

## Hardware requirements

SME requires a minimum of one MDS 9222i switch or one MSM-18/4 module in each cluster. The SME engines on the switch or module provide transparent encryption and compression to hosts and storage devices. A smart card reader is required to take advantage of all of the standard and advanced security levels.

## Software requirements

Table 206 (page 365) lists the software requirements for switches and modules in the SME cluster.

**Table 206 SME software requirements**

Component	Software version
Fabric Manager web client	Fabric Manager 3.2(3) (or later)
MDS switches attached to tape devices	SAN-OS 3.2(3) (or later)
MDS 9222i switches and switches that include the MSM-18/4 module	

## C-series SAN-OS security

This section describes the C-series SAN-OS security features.

### Simple Network Management Protocol

SNMP is an application-layer protocol that facilitates the exchange of management information between network devices. C-series switches support the following SNMP versions:

- **SNMP v1 and SNMP v2c**—Use a community-string match for user authentication.
- **SNMP v3**—Provides secure access to devices by using the following:
  - **Message integrity**—Ensures that a packet has not been tampered with while in transit
  - **Authentication**—Confirms that the message comes from a valid source
  - **Encryption**—Scrambles the packet contents, which prevents unauthorized viewing

### Remote Authentication Dial-In User Service

RADIUS is a distributed client-server protocol that protects networks against unauthorized access. RADIUS clients run on C-series switches and send authentication requests to a central RADIUS server, which contains all user authentication and network service information.

### Terminal Access Controller Access Control System

TACACS+ is a client-server protocol that uses TCP for transport. All C-series switches provide centralized authentication using TACACS+, which provides:

- Independent, modular AAA facilities
- Reliable transfers by using TCP to send data between the AAA client and server
- Encryption of all data between the switch and AAA server, which ensures data confidentiality (RADIUS encrypts passwords only)

### FC-SP and Diffie-Hellman CHAP

FC-SP provides switch-to-switch and host-to-switch authentication, which provides security challenges for large SAN fabrics. DHCHAP provide authentication between C-series switches and other devices.

## Port security

C-series port security features prevent unauthorized access to a switch port by:

- Rejecting login requests from unauthorized Fibre Channel devices or switches
- Reporting all intrusion attempts to the SAN administrator through system messages
- Using the CFS infrastructure for configuration, distribution, and restricting it to CFS-enabled switches

## Fabric binding

C-series switches in a fabric binding configuration ensure that ISLs are enabled between authorized switches only. This feature prevents unauthorized switches from disrupting traffic or joining the fabric. The EFMD protocol compares the list of authorized switches on each switch in the fabric.

## C-series IP SAN security

This section describes the C-series IP SAN security features.

### IPsec

C-series IPsec features ensure secure transmissions at the network layer. IPsec protects and authenticates IP packets between participating devices (peers) over unprotected networks. IPsec provides the following security services:

- **Data confidentiality**—Packets are encrypted by the sending device before transmitting them over the network.
- **Data integrity**—Packets are authenticated by the receiving device to ensure that data has not been altered during transmission.
- **Data-origin authentication**—The packet source can be authenticated by the receiving device.
- **Anti-replay protection**—Replayed packets can be detected and rejected by the IPsec receiver.

### CHAP authentication

C-series IP modules support CHAP, which uses a three-way handshake to ensure that validity of remote clients. C-series CHAP requires that you configure a password, which the switch presents to the iSCSI initiator. This password is used to calculate a CHAP response to a CHAP challenge sent to the IP port by the initiator.

## B-series Encryption Switch and Encryption FC Blade security

This section describes the security features for the B-series Encryption Switch and Encryption FC Blade. For switch models and fabric rules, see “B-series switches and fabric rules” (page 83).

The B-series Encryption Switch is a high-performance, 32-port autosensing 8 Gb/s Fibre Channel switch with data encryption/decryption and data compression capabilities. The switch is a network-based solution that secures data-at-rest for tape and disk array LUNs using IEEE standard AES 256-bit algorithms. Encryption and decryption engines provide in-line encryption services with up to 96 Gb/s throughput for disk I/O (mix of ciphertext and cleartext traffic).

For more information about the B-series Encryption Switch, including deployment scenarios, see the *Fabric OS Encryption Administrator's Guide*, available at [http://h18006.www1.hp.com/storage/saninfrastructure/switches/encrypt\\_sanswitch.html](http://h18006.www1.hp.com/storage/saninfrastructure/switches/encrypt_sanswitch.html).

### Features

- High-performance, scalable fabric-based encryption to enforce data confidentiality and privacy requirements
- Unparalleled encryption processing at up to 96 Gb/s to support heterogeneous enterprise data centers

- Integration with HP Secure Key Manager, providing secure and automated key sharing between multiple sites to ensure transparent access to encrypted data
- Industry-standard AES 256-bit encryption algorithms for disk arrays on a single security platform for SAN environments
- Frame Redirection technology that enables easy, nonintrusive deployment of fabric-based security services
- Plug-in encryption services available to all heterogeneous servers, including virtual machines, in data center fabrics
- Scalable performance with on-demand encryption processing power to meet regulatory mandates for protecting data

### Hardware requirements

You can use either the Encryption SAN Switch or the Encryption FC Blade for data encryption as part of the B-series Encryption Switch security platform.

### Supported security components

B-series Encryption Switch security platform supports the following software components:

- Encryption
- Frame filtering
- Advanced Zoning
- WebTools
- Enhanced Group Management

The B-series Encryption Switch security platform supports the following optional software components:

- Encryption SAN Switch Power Pack+ Software Bundle (optional)
- Adaptive Networking
- Fabric Watch
- Advanced Performance Monitor
- Extended Fabrics
- ISL Trunking
- Integrated Routing
- Data Center Fabric Manager Enterprise

## B-series Fabric OS security

This section describes the B-series Fabric OS security features for resource protection, data protection, and security validation.

### Resource protection

This section describes the B-series Fabric OS resource protection features.

#### User management

Fabric OS provides two options for authenticating users:

- **Remote RADIUS services**—Users are managed by a remote RADIUS server. All switches in the fabric can be configured to authenticate against this centralized database.
- **Local user database**—Users are managed by a local database, which is synchronized manually using the `distribute` command. This command pushes a copy of the switch's database to all other Fabric OS 5.3.0 (or later) switches in the fabric.

Fabric OS uses RBAC to determine which commands are supported for each user.

### Secure Shell

Fabric OS supports SSH encrypted sessions to ensure security. SSH encrypts all messages, including client transmission of passwords during login. SSH includes a daemon (`sshd`), which runs on the switch and supports many encryption algorithms, such as Blowfish-CBC and AES.

Commands that require a secure login channel must be issued from an original SSH session. Nested SSH sessions will reject commands that require a secure channel.

---

**NOTE:** Fabric OS 4.1.0 (or later) supports SSH V2.0 (`ssh2`).

To ensure a secure network, avoid using Telnet or any other unprotected applications to communicate with switches.

---

### Hypertext Transfer Protocol over SSL

B-series WebTools support the use of HTTPS.

The SSL protocol provides secure access to a fabric through web-based management tools like B-series WebTools. Switches configured for SSL grant access to the management tools through HTTPS links. SSL uses PKI encryption to protect data. PKI is based on digital certificates obtained from an Internet CA, which acts as the trusted agent. These certificates are based on the switch IP address or fully qualified domain names.

---

**NOTE:** If you change the switch IP address or domain name after activating its digital certificate, you may need to obtain and install a new certificate.

---

### Browser and Java support

Fabric OS 4.4.0 (or later) supports the following browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Sun Solaris and Red Hat Linux)

---

**NOTE:** In countries that allow the use of 128-bit encryption, use the current version of the browser. Upgrade to the Java 1.5.0\_06 plug-in on the management station.

---

### SNMP

B-series switches have an SNMP agent and MIB, which allow the administrator to program tools to set up switch variables and enterprise-level management processes. The SNMP ACL allows the administrator to restrict SNMP get and set operations to particular hosts and IP addresses, which provides enhanced security for the SAN.

---

**NOTE:** B-series switches support SNMP v3 and SNMP v1.

---

### Secure Copy

SCP uses SSH to securely transfer files between systems. The administrator can set the Fabric OS `configure` command to use SCP for uploads and downloads.

---

**NOTE:** FTP is not a secure protocol. File contents are in clear text during transfer, including remote login information. This limitation affects the following commands: `saveCore`, `configUpload`, `configDownload`, and `firmwareDownload`.

---



## IPFilter policy

The B-series IPFilter policy applies a set of rules to IP management interfaces as a packet filtering firewall. The firewall permits or denies traffic through the IP management interfaces according to policy rules.

Consider the following when setting IPFilter policies:

- Fabric OS supports multiple IPFilter policies, which can be defined at the same time. Each policy is identified by name and has an associated IPFilter type (IPv4 or IPv6). Do not mix IPFilter and IP address types. You can have up to six IPFilter policies defined, but only one IPFilter policy for each IPFilter type can be activated on the management IP interface.
- Audit messages are generated for changes to the IPFilter policies.
- The IPFilter policy rules are examined one by one in a list until the end of the list is reached.
- To ensure optimal performance, the most important rules should be listed first.

## Data protection

This section describes features for data protection with B-series Fabric OS.

### Fibre Channel ACLs

B-series Fabric OS uses ACLs to restrict access to data resources based on defined policies.

Fabric OS provides the following policies:

- **FCS policy**—Determines which switches can change fabric configurations
- **DCC policies**—Determines which Fibre Channel device ports can connect to which switch ports
- **SCC policy**—Determines which switches can join with another switch
- **IPFilter policy**—Filters traffic based on IP addresses

Each supported policy is identified by name; only one policy of each type can exist (except for DCC policies).

Table 207 (page 369) describes the methods for identifying policy numbers.

**Table 207 Methods for identifying policy numbers**

Policy	Device port WWN	Switch port WWN	Domain ID	Switch name
FCS_POLICY	No	Yes	Yes	Yes
DCC_POLICY_nnn	Yes	Yes	Yes	Yes
SCC_POLICY	No	Yes	Yes	Yes

## Authentication policy

By default, Fabric OS uses DHCHAP or FCAP for switch authentication. These protocols use shared secrets and digital certificates, based on switch WWN and PKI technology. Authentication automatically defaults to FCAP if both switches are configured for FCAP.

Consider the following when configuring authentication with Fabric OS:

- Fabric OS 5.3.0 (or later) is required for DHCHAP.
- DHCHAP requires the definition of a pair of shared secrets, known as a secret key pair. Each switch can share a secret key pair with any other switch or host in the fabric.
- PKI certificates must be installed on both switches to use FCAP.
- DHCHAP and FCAP are not compatible with SLAP, which is the only protocol supported in Fabric OS 3.1 and 4.2.
- Fabric OS 5.3.0 switch-to-switch authentication is backward compatible with 3.2, 4.2, 4.4, 5.0, 5.1, and 5.2.

- In the default configuration, FCAP authentication is tried first, then DHCHAP authentication. Each switch can be configured to negotiate one or both types.
- The Authentication policy is designed to accommodate mixed fabric environments that include switches running Fabric OS 5.3.0 (and earlier).
- When the Authorization policy is activated, you cannot implement a B-series Secure Fabric OS environment.

### E\_Port Authentication

The E\_Port Authentication policy allows you to configure DHCHAP authentication on the switch. By default, the policy is set to PASSIVE.

### Device Authentication policy

The Device Authentication policy is specific to HBAs. Fabric-wide distribution of the Device Authentication policy is not supported because:

- You must set the HBA and switch shared secrets manually.
- Most HBAs do not support the defined DH groups used in DHCHAP.

---

**NOTE:** By default, the switches are set to OFF, causing the security bit to be cleared during fabric login.

---

### Zones

For detailed information about B-series switch zoning, see [“Zoning guidelines for B-series Fibre Channel switches”](#) (page 113).

### B-series IP SAN

B-series IPsec uses cryptographic security to ensure private, secure communications over IP networks. Consider the following when using IPsec with B-series switches:

- IPsec is disabled by default when creating FCIP tunnels.
- IPsec provides greater security with tunneling on the B-series MP Router Blade or MP Router. IPsec does not require that you configure security for each application that uses TCP/IP. When configuring IPsec, you must ensure that either an MP Router Blade or MP Router is at each end of the FCIP tunnel.
- IPsec supports FCIP tunnels with or without IP compression, FCIP fastwrite, or tape pipelining.

### B-series iSCSI Blade

B-series iSCSI Blade supports CHAP authentication for iSCSI initiator authentication.

## Security validation

B-series Fabric OS supports a logging mechanism that captures and tracks events that are vital to security validation.

## Key management

HP StorageWorks Secure Key Manager for HP LT04 tape libraries is part of the Secure Advantage solution.

Secure Key Manager features include:

- Centralized encryption key management for HP LT04 tape libraries
  - Automatic policy-based key generation and management supporting key and cartridge

- granularity
- ISV transparent key archival and retrieval for multiple libraries
- Extensible to emerging open standards
- Strong auditable security for encryption keys to ensure compliance
  - Hardened server appliance
  - Secure identity-based access, administration, and logging
  - Designed for FIPS 140-2 validation
- Reliable lifetime key archival, which ensures key availability, even in the event of a site disaster
  - Automatic multiple-site clustering, key encryption, and failover
  - Comprehensive backup and restore functions for keys
  - Redundant device components and active alerts

### Secure Key Manager validation process

The HP CLW enhances the Secure Key Manager validation process.

The CLW features include:

- High-performance appliance with the following modules: Log Manager, Analysis Manager, and Real-time Alert Manager
- High-speed collection and analysis of log data, which automates compliance reporting for industry and government standards

### Integration of key management with partners

Secure Key Manager can be integrated with third-party and partner products (such as the C-series SME) to provide a standard enterprise data security solution.

---

## 21 Best practices

This chapter describes HP best practices for SAN design and implementation. It describes the following topics:

- “SAN planning” (page 372)
- “Design specification” (page 372)
- “SAN topology” (page 373)
- “SAN configuration ” (page 374)
- “Storage-based LUN masking” (page 375)
- “Zoning” (page 376)
- “FCoE switch configuration quick-setup instructions” (page 379)
- “SAN scaling” (page 389)
- “SAN fabric merging” (page 390)

For SAN design assistance, see the HP Storage Services website <http://www8.hp.com/us/en/services/services-detail.html?compURI=tcm:245-808667&pageTitle=Storage-Technology-Services>.

### SAN planning

Allocate adequate time to plan your SAN prior to implementation. Design a SAN that fulfills current and future requirements for capacity and connectivity.

During the planning phase, consider these design factors and recommendations:

- **Deployment strategy**  
Consider initially implementing entry-level SANs that can be interconnected to increase capacity. Entry-level SANs are relatively easy to implement. Enterprise SANs offer economy of scale; however, they are more complex and can take longer to implement.
- **Topology design**  
Choose an initial design that can accommodate expansion without transitioning to a different topology.
- **Experience level**  
If you have limited experience implementing a SAN, start with an entry-level SAN. As you gain experience, deploy mid-range or enterprise-level SANs.
- **SAN management strategy**  
Specify the policies, identification schemes, and tools to manage your SAN.
- **Technological advances**  
Anticipate the gradual availability of Fibre Channel switches that have more ports and faster interconnect speeds.

### Design specification

During the planning process, create a specification that describes your decisions and design. Review and evaluate the design, compare alternatives, make adjustments, and communicate plans before implementation.

A complete design specification includes the following elements:

- **Topology map**—Shows the logical SAN topology and fabric interconnect scheme; describes a strategy to accommodate expansion and technological advances
- **Configuration layout**—Shows the physical layout of components; use for troubleshooting and to verify the correct connectivity
- **Storage map**—Defines the storage system configuration and settings, such as host LUN allocation and RAID levels
- **Zoning map**—Defines the communication access settings for devices and user ports in the SAN

## SAN topology

This section describes SAN features for enterprise-level SANs:

- [“Multi-fabric SANs” \(page 373\)](#)
- [“Failover protection” \(page 373\)](#)
- [“Data access patterns” \(page 374\)](#)
- [“ISL ratio” \(page 374\)](#)
- [“Incremental SAN expansion” \(page 374\)](#)

## Multi-fabric SANs

HP recommends using two or more separate fabrics for enterprise-level SANs. Multiple fabrics protect against potential failure points, such as hardware, software, or operator error. The failure of one fabric does not affect other fabrics in the SAN.

## SAN and fabric monitoring

For all single and multiple fabric configurations, HP recommends that you utilize intelligent fabric monitoring tools such as B-series Fabric Watch and C-series RMON (see the user documentation for your switch). These products provide detailed monitoring of individual Fibre Channel ports, notifications, and in some cases, automated isolation of inoperative devices. This can help you to identify abnormal conditions and avoid operational degradation. This degradation can adversely affect operation, but does not necessarily result in a failover event.

---

**NOTE:** These tools are effective in detecting and avoiding most abnormal conditions. Failing or marginal hardware can cause rare conditions which are not detected. In order to minimize these rare conditions, HP recommends that you implement proper cable management practices and end-to-end SAN monitoring.

---

## Failover protection

Use failover technology in SAN configurations that have two or more fabrics. Each server has two or more HBAs. If the communication path from one HBA to the storage system fails, the I/O traffic is rerouted through the other HBA.

To minimize the risk of uneven workloads, configure the separate fabrics for similar size and topology. You can also use failover protection in SANs with only one fabric to protect against HBA, path, and storage controller failures. For more information, see [“Data availability” \(page 38\)](#).

## Data access patterns

Review your data access needs before making a topology choice. The optimum SAN configuration depends on I/O traffic requirements and data access patterns:

- **Local (one-to-one)**—Data access between a local server and a storage system connected to the same switch
- **Centralized (many-to-one)**—Data access between multiple, dispersed servers and one centrally located storage system
- **Distributed (many-to-many)**—Data access between multiple, dispersed servers and storage systems

For recommended topologies based in the primary data access type, see [“Data access performance by SAN fabric topology”](#) (page 33).

## ISL ratio

Determine the ISL ratio for switch-to-switch connectivity based on the workload of servers and storage systems. In some cases, you can assess the I/O requirements of your applications and servers by using application sizing tools. After deployment, use current system measurements to determine the actual workload, and modify your implementation if the initial design does not meet your requirements. For more information, see [“Recommended ISL ratios”](#) (page 31).

## Incremental SAN expansion

Plan for expansion when the initial design is a subset of a future, larger design. For example, if you are using a core-edge topology in the initial design, allocate spare ports on the core switches to support addition of edge switches.

To expand the SAN, you can make incremental changes rather than reconfiguring the entire SAN. Changes to the core switches are isolated from the edge switches, which minimizes the effect of changes required to support core growth. Changes to the server connection to an edge switch are isolated from the core, which minimizes the effect of server-related changes. If you are using two or more fabrics, you can temporarily route server I/O traffic to one fabric while the other fabric is being modified.

## SAN configuration

After completing the planning phase, you can configure your SAN. During the configuration phase, record details about the physical configuration.

To facilitate maintenance, observe the following practices:

- **Record keeping**  
Record the cable connections on the configuration layout diagram. Record the WWN and location of each node and device.
- **HBA labels**  
Affix a label on each HBA that clearly identifies the WWN. HP storage systems are prelabeled with this information. Affix another label in plain view if necessary.
- **Cable labels**  
Label both ends of each cable with a cable number or color-coding scheme. This allows you to quickly identify each cable. Securely affix a label to each end of the cable to identify connection points, such as TO and FROM.
- **Switch ports**  
Use port plugs to protect unused switch ports; never leave ports exposed.

- **Cable dressing**

Use care when routing fiber optic cable and ensure that cables conform to the minimum bend radius requirements. See [“Rules for fiber optic cable connections ” \(page 147\)](#). Use hook-and-loop (such as Velcro brand) tie wraps to group and support the cables.

---

△ **CAUTION:** Plastic tie wraps can damage the internal fiber core if over-tightened.

---

- **Cable symmetry.**

When connecting cables, use similar slot and port numbers. For example, connect HBA 1 to SAN fabric 1, HBA 2 to SAN fabric 2, and so on.

Consider the following when configuring your SAN:

- [“Fibre Channel switch configuration” \(page 375\)](#)
- [“Server setup” \(page 375\)](#)
- [“Storage system configuration” \(page 375\)](#)

## Fibre Channel switch configuration

HP Fibre Channel switches are preconfigured with compatible parameter settings. You must configure each switch with a unique domain ID.

## Server setup

When setting up servers:

- For each platform or operating system type, use only HP-supported drivers and configuration settings.
- Ensure that the servers have the supported operating system versions and all required updates.
- Use an alphanumeric naming scheme for multiple servers of the same type, such as WIN01 and WIN02 for Windows servers.

## Storage system configuration

When configuring storage systems:

- Use the storage map created in the planning phase to configure each storage system.
- Verify server-to-storage connectivity, and access one server at a time.
- When defining storagesets, disable all access first, and then enable the desired access.
- Define storage system connection names that are consistent with zoning alias names. Use consistent names for storage port and controller connections. Use a naming scheme that represents the physical connectivity.

## Storage-based LUN masking

Use storage-based LUN masking to allow or prevent access to storage system LUNs from one or more servers.

- For HP EVA storage systems, use SSP.
- For HP XP storage systems, use LUN Configuration and Security Manager XP.
- For HP VA storage systems, use Secure Manager VA.

During SCSI initialization, the server queries the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list and returns the LUNs assigned to the WWN. Any other LUNs on that storage port are not available to the server.

## Zoning

This section describes configuration recommendations for:

- [“Zoning enforcement” \(page 376\)](#)
- [“Zoning guidelines” \(page 376\)](#)
- [“EBS zoning” \(page 377\)](#)
- [“Zone naming” \(page 377\)](#)

### Zoning enforcement

To protect against unauthorized access, Fibre Channel switches provide three types of zoning enforcement (listed here in order of enforcement):

- **Access authorization**

Access authorization provides frame-level access control in hardware and verifies the SID-DID combination of each frame. The frame is delivered to the destination only if specified as a valid combination in the zone definition. This method offers a high level of security and is classified as hard zoning because it requires hardware resources at the ASIC level.
- **Discovery authentication**

Discovery authentication occurs during access to the NS directory. The fabric presents only a partial list of authorized devices from the NS directory. This method may be enforced by software or hardware, depending on the switch model. When enforced by software, this method is susceptible to security threats from unauthorized devices that violate Fibre Channel protocols.
- **Soft-plus zoning by login authentication**

In addition to discovery authentication, some switches enforce authentication at the Fibre Channel protocol login frame level. For example, if a host sends a PLOGI frame to a device that is not a member of its zone, the frame is dropped. Login authentication provides more protection than discovery authentication but is not as secure as access authorization.

The zone configuration and the switch model determine the type of zoning enforcement you can implement in your SAN fabric. For information about the relationship of zone configuration with zoning enforcement, see the following tables:

- [Table 37 \(page 113\)](#) (B-series)
- [Table 52 \(page 129\)](#) (C-series)
- [Table 60 \(page 139\)](#) (H-series)
- [Table 70 \(page 145\)](#) (M-series)

Some system restrictions affect the movement of devices within the fabric, regardless of zoning type. For example, some operating systems, such as HP-UX, create device file names based on the 24-bit fabric address and do not allow moving the device to a different port. A change in the address causes the device to be treated as a different device.

### Zoning guidelines

Use one of the following zoning methods:

- Operating system (minimum level required)
- HBA
- Application
- Port allocation



## Zoning by operating system

Zoning by operating system is the minimal required zoning method. This method allows multiple HBAs with the same operating system to be grouped with the accessed storage ports. Zoning by operating system prevents the interaction of HBAs with incompatible operating systems.

This method limits the number of zones in a fabric. A large zone can be divided into multiple zones within the operating system type. Zoning by operating system type limits disruptions and the number of fabric change notifications.

Certain situations require zoning by HBA, for example, configuring server access to multiple storage types. For additional zoning requirements, see:

- “Common server access” (page 201)
- “Zoning limits and enforcement” (page 138) (H-series)
- “Zoning limits and enforcement” (page 112) (B-series)
- “Zoning limits and enforcement” (page 129) (C-series)
- “Zoning limits and enforcement” (page 144) (M-series)

## Zoning by HBA

For zoning by HBA, each zone has only one HBA (initiator); each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, HP recommends that a second zone be created for the HBA and associated tape devices. For zoning requirements with different HBA models on the same server, see “Common server, different HBAs” (page 201).

This zoning philosophy is the preferred method for both standalone and clustered systems; zoning by single HBA requires the creation of numerous zones; each containing only a few members. Zone changes affect a small number of devices, minimizing the effect of an incorrect zone change.

## Zoning by application

Zoning by application configures multiple, sometimes incompatible, operating systems into the same zones. This method allows the potential for disruptions among servers, such as a web server disrupting a data warehouse server. A zone with a large number of members is susceptible to more administrative errors, such as distribution of RSCNs to a larger group than necessary.

## Zoning by port allocation

Avoid zoning by port allocation unless you have strictly enforced processes for port and device allocation in the fabric.

There is no consequence for the change of WWN when a storage port, server HBA, or tape drive is replaced. If the new device connects to the original port, it continues to have the same access rights. You can preassociate switch ports with storage ports, and therefore control the server-to-storage ratio. This technique prevents overloading of a storage port by allowing you to limit the number of servers that are allowed access.

## EBS zoning

For EBS zoning recommendations, see the *HP StorageWorks Enterprise Backup Solution Design Guide*, available at <http://h18004.www1.hp.com/products/storageworks/ebs/documentation.html>.

## Zone naming

When naming zones:

- Configure and test small zones that are a subset of the larger SAN.
- Use meaningful names for zones and zone member aliases.

- Use a consistent naming scheme for all components.
- Before making zoning changes, save the current configuration.
- If possible, do not make zoning changes when a switch in the fabric is temporarily unavailable.

## Naming by identifier type

Define zone member names by using one of the following naming conventions:

- **Domain ID and port number**  
Use the switch domain ID and port number to identify zone members. The zone definition remains intact when an HBA or controller is replaced by another with a different WWN. However, when a device is moved to a different port in the fabric, it is no longer a member of the zone.
- **WWN**  
Use the device WWN to uniquely identify zone members. The zone definition remains intact when the device is moved to a different port or switch in the fabric. However, if an HBA is replaced by another with a different WWN, you must update the zone definition.
- **WWN with domain ID and port number**  
Use the switch domain ID, port number, and WWN to identify zone members.

## Case sensitivity of fabric identifiers

To define an alias naming scheme, consider the case sensitivity of fabric identifiers:

- **Case sensitive**—Switch alias names
- **Not case sensitive**—Device connection names

### Case sensitivity example

RING\_1 and Ring\_1 are distinct switch identifiers.

## Server naming

Servers are identified by the WWN of the HBA. For server aliases, use the operating system name and the HBA number. For example, for server WIN01 with one HBA, define the alias as WIN01\_HBA01; for the second HBA, define the alias as WIN01\_HBA02.

## Storage system naming

Fibre Channel storage systems have a unique WWN for each controller port. When implementing multiple fabrics, different ports are configured in each fabric.

Observe the following best practices:

- Assign a unique number to each storage port.  
For example, ports in fabric 1 have aliases S1\_A1 and S1\_A2.
- Use a consistent alias naming convention for ports and HBAs throughout the configuration.  
For example, configure HBA 1, S1\_A1 and S1\_B2 in fabric 1; configure HBA 2, S1\_A2 and S1\_B1 in fabric 2.
- Define host connection names for the HBA WWNs similar to the alias names in the fabric.  
For example, define alias WIN01\_HBA1 for Windows server 1 and HBA 1.

### Alias convention example

A storage system connected to two fabrics has the following aliases:

- **First fabric:**
  - S1\_A1
  - S1\_B2
- **Second fabric:**
  - S1\_A2
  - S1\_B1

Ports A1 and B2 are cabled to the first fabric. Ports A2 and B1 are cabled to the second fabric.

## FCoE switch configuration quick-setup instructions

FCoE CN switches have dual capabilities in that they serve as both an Ethernet switch and an FC switch. You must perform a setup procedure to achieve the desired function.

- 
- ❗ **IMPORTANT:** The procedures in this section are intended as a quick-start guide to configuring the FCoE ports on an FCoE CN switch to access the FC ports on the FCoE CN switch and other attached FC switches. It is intended for users who are familiar with FCoE switches and their associated configuration commands.
- 

### HP A5820 FCoE Converged Network Switch quick setup

This procedure is intended for users who are familiar with HP A-series switches. Use this procedure to enable servers with CNAs attached to the HP A5820 FCoE Converged Network Switch to access devices from an attached B-series, C-series, or H-series FC fabric.

- 
- ❗ **IMPORTANT:** If you are not familiar with HP A-series switches, refer to the A5820 product documentation for detailed instructions on setting up the switch.
- 

1. Create 3 VLANs (for LAN, SAN, and SAN discovery traffic).

```
<switchname> system-view
[switchname] vlan 1001
[switchname-vlan1001] description ToLAN
[switchname-vlan1001] quit
[switchname] vlan 4001
[switchname-vlan4001] description ToSAN
[switchname-vlan4001] quit
[switchname] vlan 3001
[switchname-vlan3001] description FIPVLAN
[switchname-vlan3001] protocol-vlan 0 mode ethernetii etype 8914
[switchname-vlan3001] quit
[switchname]
```

2. Configure DCBX.

```
[switchname] acl number 4000 name DCBX
[switchname-acl-ethernetframe-4000-DCBX] rule 0 permit type 8906 ffff
[switchname-acl-ethernetframe-4000-DCBX] rule 5 permit type 8914 ffff
[switchname-acl-ethernetframe-4000-DCBX] quit
[switchname] traffic classifier DCBX operator or
[switchname-classifier-DCBX] if-match acl 4000
[switchname-classifier-DCBX] quit
[switchname] traffic behavior DCBX
[switchname-behavior-DCBX] remark dot1p 3
[switchname-behavior-DCBX] quit
[switchname] qos policy DCBX
[switchname-qospolicy-DCBX] classifier DCBX behavior DCBX mode dcbx
[switchname-qospolicy-DCBX] quit
```

3. Configure ETS.

```

[switchname] qos map-table dot1p-lp
[switchname-maptbl-dot1p-lp] import 0 export 0
[switchname-maptbl-dot1p-lp] import 1 export 0
[switchname-maptbl-dot1p-lp] import 2 export 0
[switchname-maptbl-dot1p-lp] import 3 export 1
[switchname-maptbl-dot1p-lp] import 4 export 0
[switchname-maptbl-dot1p-lp] import 5 export 0
[switchname-maptbl-dot1p-lp] import 6 export 0
[switchname-maptbl-dot1p-lp] import 7 export 0
[switchname-maptbl-dot1p-lp] quit

```

#### 4. Configure Interfaces (3 types).

- Server CNA Interfaces

```

[switchname] interface Ten-GigabitEthernet1/0/1
[switchname-interface Ten-GigabitEthernet1/0/1] port link-type hybrid
[switchname-interface Ten-GigabitEthernet1/0/1] port hybrid vlan 4001 tagged
[switchname-interface Ten-GigabitEthernet1/0/1] port hybrid vlan 1001 3001 untagged
[switchname-interface Ten-GigabitEthernet1/0/1] port hybrid pvid vlan 1001
[switchname-interface Ten-GigabitEthernet1/0/1] undo port hybrid vlan 1
[switchname-interface Ten-GigabitEthernet1/0/1] port hybrid protocol-vlan vlan 3001 0
[switchname-interface Ten-GigabitEthernet1/0/1] qos apply policy DCBX outbound
[switchname-interface Ten-GigabitEthernet1/0/1] lldp tlv-enable dot1-tlv dcbx
[switchname-interface Ten-GigabitEthernet1/0/1] priority-flow-control auto
[switchname-interface Ten-GigabitEthernet1/0/1] priority-flow-control no-drop dot1p 3
[switchname-interface Ten-GigabitEthernet1/0/1] qos trust dot1p
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 7 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 6 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 5 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 4 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 3 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 2 group sp
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 1 group byte-count 15
[switchname-interface Ten-GigabitEthernet1/0/1] qos wrr 0 group byte-count 15
[switchname-interface Ten-GigabitEthernet1/0/1] broadcast-suppression 1
[switchname-interface Ten-GigabitEthernet1/0/1] multicast-suppression 1
[switchname-interface Ten-GigabitEthernet1/0/1] unicast-suppression 1
[switchname-interface Ten-GigabitEthernet1/0/1] quit

```

- Network (No SAN Access) Interfaces

```

[switchname] interface Ten-GigabitEthernet1/0/14
[switchname-interface Ten-GigabitEthernet1/0/14] port link-type trunk
[switchname-interface Ten-GigabitEthernet1/0/14] port trunk permit vlan 1001
[switchname-interface Ten-GigabitEthernet1/0/14] port trunk pvid vlan 1001
[switchname-interface Ten-GigabitEthernet1/0/14] undo port trunk permit vlan 1
[switchname-interface Ten-GigabitEthernet1/0/14] quit

```

- Internal A5820X FCoE module Interfaces (1/1/1-1/1/4 and/or 1/2/1-1/2/4)

```

[switchname] interface Ten-GigabitEthernet1/1/1
[switchname-interface Ten-GigabitEthernet1/1/1] port link-type hybrid
[switchname-interface Ten-GigabitEthernet1/1/1] port hybrid vlan 4001 tagged
[switchname-interface Ten-GigabitEthernet1/1/1] port hybrid vlan 3001 untagged
[switchname-interface Ten-GigabitEthernet1/1/1] port hybrid pvid vlan 3001
[switchname-interface Ten-GigabitEthernet1/1/1] undo port hybrid vlan 1
[switchname-interface Ten-GigabitEthernet1/1/1] stp disable
[switchname-interface Ten-GigabitEthernet1/1/1] lldp tlv-enable dot1-tlv dcbx
[switchname-interface Ten-GigabitEthernet1/1/1] priority-flow-control auto
[switchname-interface Ten-GigabitEthernet1/1/1] priority-flow-control no-drop dot1p 3
[switchname-interface Ten-GigabitEthernet1/1/1] qos trust dot1p
[switchname-interface Ten-GigabitEthernet1/1/1] quit

```

#### 5. Save changes and exit system-view.

```

[switchname] save
[switchname] quit
<switchname>

```

6. Configure A5820 FCoE module.
  - a. Connect to FCoE module from A5820 CLI.
 

```
<switchname>oap connect slot 1 system subslot1
```

At the Login prompt, type the initial default user account, **admin**. At the Password prompt, type the initial default password, **password**.

```
H3C #>admin start
H3C (admin) #>config edit
```
  - b. Create VLAN to match A5820 ToSAN VLAN.
 

```
H3C (admin-config) #> vlan 4001 create
H3C (admin-config) #> vlan 4001 add port 4-7
```
  - c. Add ToSAN VLAN to default FCF.
 

```
H3C (admin-config) #> fcf list
```

Verify default FCF and VLAN are FCF: 0EFC00 VLAN: 1002.

```
H3C (admin-config) #> fcf 0EFC00 remove vlan 1002
H3C (admin-config) #> fcf 0EFC00 add vlan 4001
```
  - d. Save and activate changes.
 

```
H3C (admin-config) #>config save
H3C (admin) #>config activate
H3C (admin) #>admin stop
```
  - e. Optionally configure FCoE module IP address for management access.
 

```
H3C #> admin start
H3C #>set setup system
```

(and follow prompts)
  - f. Return back to A5820 prompt.
 

```
H3C #> quit
```

Press **CTRL+K** to return to A5820 prompt.

```
<switchname>
```

## 2408 FCoE Converged Network Switch and DC SAN Director Switch 10/24 FCoE Blade quick setup

This procedure is intended for users who are familiar with Brocade FC switches and have experience merging B-series FC switches into an existing FC fabric. Use this procedure to enable servers with CNAs attached to the 2408 FCoE Converged Network Switch or the DC SAN Director Switch 10/24 FCoE Blade to access devices on the attached B-series FC fabric.

- 
- ❗ **IMPORTANT:** If you are not familiar with Brocade FC switches or you do not have experience merging B-series FC switches into an existing FC fabric, use the detailed instructions found in the switch user guide to set up your switch.
- 

1. Configure LLDP for FCoE (common to all CEE ports).
 

```
switch:admin> cmsh
switch# enable
switch# config terminal
switch(config)# protocol lldp
switch(conf-lldp)# advertise dcbx-fcoe-app-tlv
switch(conf-lldp)# advertise dcbx-fcoe-logical-link-tlv
switch(conf-lldp)# exit
```
2. Create a CEE map to carry LAN traffic (60%) and SAN traffic (40%) (common to all CEE ports).

```

switch(config)# cee-map default
switch(conf-ceemap)# priority-group-table 1 weight 40 pfc
switch(conf-ceemap)# priority-group-table 2 weight 60
switch(conf-ceemap)# priority-table 2 2 2 1 2 2 2 2
switch(conf-ceemap)# exit

```

3. Create an FCoE VLAN for traffic to and from the FC fabric (required for FCoE).

```

switch(config)# vlan classifier rule 1 proto fcoe encap ethv2
switch(config)# vlan classifier rule 2 proto fip encap ethv2
switch(config)# vlan classifier group 1 add rule 1
switch(config)# vlan classifier group 1 add rule 2
switch(config)# interface vlan 5 (Can be any VLAN number other than 1)
switch(conf-if-vl-5)# fcf forward
switch(conf-if-vl-5)# exit

```

4. Configure interfaces (required for each port being configured).

```

switch(config)# interface tengigabitethernet 0/0 Change Port ID
switch(config-if-te-0/0)# switchport
switch(config-if-te-0/0)# switchport mode converged1
switch(config-if-te-0/0)# switchport converged allowed vlan add 51 FCoE Required (Step 3 VLAN)
switch(config-if-te-0/0)# vlan classifier activate group 1 vlan 51
switch(config-if-te-0/0)# cee default1
switch(config-if-te-0/0)# exit

```

<sup>1</sup>This command allows the port to access the FCoE VLAN. You can omit this command for non-FCoE ports; however, both FCoE and non-FCoE ports might require a similar command for access to other VLANs.

5. Repeat [Step 4](#) for each interface you are configuring.
6. Perform a port shutdown/no shutdown operation (required for each port for the configuration change to be enabled).

```

switch(config)# interface tengigabitethernet 0/0 Change Port ID
switch(config-if-te-0/0)# shutdown
switch(config-if-te-0/0)# no shutdown
switch(config-if-te-0/0)# exit

```

7. Repeat [Step 6](#) for each interface you are configuring.
8. Save the running configuration to boot flash.

```

switch(config)# exit
switch# copy running-config startup-config
Overwrite the startup config file (y/n): y
Building configuration...
switch#

```

9. Verify that the CEE port link status and VLAN status are correct.

```

switch# show ip interface brief
Interface                IP-Address      Status          Protocol
=====
TenGigabitEthernet 0/0  unassigned    up              up

switch#show vlan brief

VLAN    Name            State  Ports
=====
1       default        ACTIVE Te 0/0 (u)...
5       VLAN0005       ACTIVE Te 0/0 (u)...

switch# exit
switch:admin>

```

## 10. Verify the status of the FC and FCoE virtual FC ports.

```
BR8000-01:admin> switchshow
switchName: BR8000-1
switchType: 76.7
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 4
switchId: fffc04
switchWwn: 10:00:00:05:1e:76:77:80
zoning: ON (Brocade_East)
switchBeacon: OFF
FC Router: OFF
FC Router BB Fabric ID: 1
```

Index	Port	Address	Media	Speed	State	Proto			
0	0	040000	id	N4	Online	FC	E-Port	10:00:00:05:1e:36:2a:70	"BR48-02"
1	1	040100	id	N4	Online	FC	E-Port	10:00:00:05:1e:36:2a:70	"BR48-02"
2	2	040200	id	N4	Online	FC	E-Port	10:00:00:05:1e:36:2a:70	"BR48-02"
3	3	040300	id	N4	Online	FC	E-Port	10:00:00:05:1e:36:2a:70	"BR48-02" (upstream)
4	4	040400	id	N8	No_Light	FC			
5	5	040500	id	N8	No_Light	FC			
6	6	040600	id	N8	No_Light	FC			
7	7	040700	id	N8	No_Light	FC			
8	8	040800	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
9	9	040900	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
10	10	040a00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
11	11	040b00	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
12	12	040c00	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
13	13	040d00	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
14	14	040e00	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
15	15	040f00	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
16	16	041000	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
17	17	041100	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
18	18	041200	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
19	19	041300	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
20	20	041400	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
21	21	041500	--	10G	Online	FCoE	VF-Port	1	VN-Port (s)
22	22	041600	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
23	23	041700	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
24	24	041800	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
25	25	041900	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
26	26	041a00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
27	27	041b00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
28	28	041c00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
29	29	041d00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
30	30	041e00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)
31	31	041f00	--	10G	Online	FCoE	VF-Port	0	VN-Port (s)

**NOTE:** Ports 8 through 31 are the virtual FC ports. To display the FCoE virtual FC devices connected to those ports, enter the following command.

```
BR8000-01:admin> fcoe --loginshow
=====
Port    Te port    Device WWN                Device MAC                Session MAC
=====
8       Te 0/0     10:00:00:00:c9:93:c9:13   00:00:c9:93:c9:13       0e:fc:00:04:08:01
9       Te 0/1     10:00:00:00:c9:93:c9:2b   00:00:c9:93:c9:2b       0e:fc:00:04:09:01
11      Te 0/3     10:00:00:00:c9:93:c8:b3   00:00:c9:93:c8:b3       0e:fc:00:04:0b:01
12      Te 0/4     10:00:00:00:c9:93:ca:8b   00:00:c9:93:ca:8b       0e:fc:00:04:0c:01
13      Te 0/5     10:00:00:00:c9:93:c9:bf   00:00:c9:93:c9:bf       0e:fc:00:04:0d:01
14      Te 0/6     10:00:00:00:c9:93:ca:e3   00:00:c9:93:ca:e3       0e:fc:00:04:0e:01
15      Te 0/7     10:00:00:00:c9:93:ca:eb   00:00:c9:93:ca:eb       0e:fc:00:04:0f:01
16      Te 0/8     10:00:00:00:c9:93:c9:1b   00:00:c9:93:c9:1b       0e:fc:00:04:10:01
20      Te 0/12    10:00:00:00:c9:93:ca:93   00:00:c9:93:ca:93       0e:fc:00:04:14:01
21      Te 0/13    10:00:00:00:c9:93:c8:43   00:00:c9:93:c8:43       0e:fc:00:04:15:01
BR8000-1:admin>
```

## HP C-series Nexus 5010/20 Converged Network Switch quick setup

This procedure is intended for users who are familiar with Cisco MDS FC switches and have experience merging C-series MDS FC switches into an existing MDS FC fabric. Use this procedure to enable servers with CNAs attached to the HP C-series Nexus 5010/20 Converged Network Switch to access devices on the attached C-series FC fabric.

- ❗ **IMPORTANT:** If you are not familiar with Cisco FC switches or you do not have experience merging C-series FC switches into an existing FC fabric, use the detailed instructions found in the switch user guide to set up your switch.

HP recommends that you use the VFC port assignments listed in [Table 208 \(page 384\)](#).

**Table 208 Recommended VFC port assignments**

Model	Slot	Number of FCoE ports (maximum)	VFC port assignments (recommended)
HP C-series Nexus 5010 Converged Network Switch	1 (Standard)	20	VFC ports 1 through 20
	2 (Optional expansion module slot <sup>1</sup> )	6	VFC ports 41 through 46
HP C-series Nexus 5020 Converged Network Switch	1 (Standard)	40	VFC ports 1 through 40
	2 (Optional expansion module slot <sup>1</sup> )	6	VFC ports 41 through 46
	3 (Optional expansion module slot <sup>1</sup> )	6	VFC ports 47 through 52

<sup>1</sup> Optional expansion module slots can contain either 10-GbE ports or FC ports.

In the following examples, FCoE is enabled on a C-series Nexus 5010 Converged Network Switch, VLAN 200 is created, and the Ethernet ports are bound to a virtual SAN (VSAN 2) from an 8-port FC expansion module. The commands are the same for the HP C-series Nexus 5020 Switch.



To establish CNA connectivity and enable login to the HP C-series Nexus 5000 Converged Network Switch, configure the IEEE DCB ports as follows:

1. Enable FCoE (disabled by default).

---

**NOTE:** The C-series Nexus 5000 Converged Network Switch will require a reload (reboot).

---

```
Nexus5010# configure terminal
Nexus5010(config)# feature fcoe
Nexus5010(config)# feature npiv
Nexus5010(config)# 2009 Apr 1 12:05:06 Nexus5010 %$ VDC-1 %$ %PLATFORM-2-
FC_LICENSE_DESIRED: FCoE/FC feature will be enabled after the configuration
is saved followed by a reboot
Nexus5010(config)# exit
Nexus5010# copy running-config startup-config
[#####] 100%
Packaging and storing to flash: \
Packaging and storing to flash: |
Packaging and storing to flash: /
```

---

**NOTE:** Depending on the NX-OS version, a reload may be required when enabling features. When you execute the `feature` command, a message appears to indicate if a reload is required.

---

2. Reload the system if you are directed to do so. Otherwise, proceed to [Step 3](#).

```
Nexus5010# reload
```

```
WARNING: This command will reboot the system
Do you want to continue? (y/n) [n] y
The system is going down for reboot NOW!
```

3. Create a new VLAN for FCoE.

By default, all ports are in VLAN 1, however, you must use a different VLAN for FCoE. In the following example, VLAN 200 is created with access to Ethernet ports 1/1 to 1/20, and VFC ports 1–20 are also created.

---

**NOTE:** In the last section of this example, 1/1-20 indicates that the commands that follow apply to multiple ports (in this case, ports 1/1 through 1/20). All 20 ports are set for switchport mode trunk, and switchport trunk allowed is set for VLANs 1 and 200.

---

```
Nexus5010# configure terminal
Nexus5010(config)# vlan 200
Nexus5010(config-vlan)# exit
Nexus5010(config)# exit
Nexus5010# show vlan brief
VLAN Name Status Ports
-----
1    default active  Eth1/1, Eth1/2, Eth1/3,
    Eth1/4, Eth1/5, Eth1/6,
    Eth1/7, Eth1/8, Eth1/9,
    Eth1/10, Eth1/11, Eth1/12,
    Eth1/13, Eth1/14, Eth1/15,
    Eth1/16, Eth1/17, Eth1/18,
    Eth1/19, Eth1/20
200  VLAN0200 active
```

```
Nexus5010# configure terminal
Nexus5010(config)# interface ethernet 1/1-20
Nexus5010(config-if-range)# switchport mode trunk
Nexus5010(config-if-range)# switchport trunk allowed vlan 1, 2001
```

```
Nexus5010(config-if-range)# interface vfc 1-20
Nexus5010(config-if-range)# exit
```

<sup>1</sup>This command allows the port to access the FCoE VLAN (VLAN 200 in this example). For non-FCoE ports, you can omit the FCoE VLAN from this command; however, both FCoE and non-FCoE ports might require access to other VLANs.

4. Create a new VSAN that includes the FC and VFC ports.

By default, all ports are in VSAN 1. HP recommends that you use a different VSAN for SAN connectivity. In this example, VSAN 2 is created and includes FC ports 2/1 through 2/8 and VFC ports 1 through 20.

---

**NOTE:** VFC ports must be FCoE ports. FC ports cannot be VFC ports.

---

```
Nexus5010# configure terminal
Nexus5010(config)# vsan database
Nexus5010(config-vsan-db)# vsan 2
Nexus5010(config-vsan-db)# vsan 2 interface fc2/1-8
Nexus5010(config-vsan-db)# vsan 2 interface vfc 1-20
Nexus5010(config-vsan-db)# exit
Nexus5010(config)# exit
Nexus5010# show vsan membership
vsan 1 interfaces:
```

```
vsan 2 interfaces:
fc2/1   fc2/2   fc2/3   fc2/4   fc2/5   fc2/6   fc2/7   fc2/8
vfc1    vfc2    vfc3    vfc4    vfc5    vfc6    vfc7    vfc8
vfc9    vfc10   vfc11   vfc12   vfc13   vfc14   vfc15   vfc16
vfc17   vfc18   vfc19   vfc20
```

```
vsan 4094(isolated_vsan) interfaces:
```

5. Associate the VLAN with the VSAN.

In this example, VLAN 200 is associated with VSAN 2.

```
Nexus5010# configure terminal
Nexus5010(config)# vlan 200
Nexus5010(config-vlan)# fcoe vsan 2
Nexus5010(config-vlan)# exit
Nexus5010(config)# exit
Nexus5010# show vlan fcoe
```

VLAN	VSAN	Status
-----	-----	-----
200	2	Operational

6. Bind each VFC port to a unique Ethernet port by issuing the following commands on each port:

```
interface vfc n
bind interface ethernet slot/port
exit
```

---

**NOTE:** Depending on your switch, up to 52 VFC ports may be available.

---

In [Example 1 “Creating and binding consecutive VFC ports”](#), VFC ports 1 through 8 are created and are bound to Ethernet ports 1/1 through 1/8, respectively. [Example 2 “Creating and binding nonconsecutive VFC ports”](#) shows a more complex configuration in which the VFC ports are not sequential.

## Example 1 Creating and binding consecutive VFC ports

---

```
Nexus5010# configure terminal
Nexus5010(config)# interface vfc 1
Nexus5010(config-if)# bind interface ethernet 1/1
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 2
Nexus5010(config-if)# bind interface ethernet 1/2
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 3
Nexus5010(config-if)# bind interface ethernet 1/3
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 4
Nexus5010(config-if)# bind interface ethernet 1/4
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 5
Nexus5010(config-if)# bind interface ethernet 1/5
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 6
Nexus5010(config-if)# bind interface ethernet 1/6
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 7
Nexus5010(config-if)# bind interface ethernet 1/7
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 8
Nexus5010(config-if)# bind interface ethernet 1/8
Nexus5010(config-if)# exit
Nexus5010(config)# interface vfc 1-8
Nexus5010(config-if)# no shutdown
Nexus5010(config-if)# exit
```

- 
- ① **IMPORTANT:** Example 2 “Creating and binding nonconsecutive VFC ports” is provided for reference only. It shows an alternate method of performing this step.
- 

In Example 2 “Creating and binding nonconsecutive VFC ports”, VFC ports 1 through 6, 10, and 20 through 25 are created, and each VFC port is bound to an Ethernet port. Because the VFC ports are not sequential, multiple `interface vfc` commands are required.

## Example 2 Creating and binding nonconsecutive VFC ports

---

```
Nexus5020# configure terminal
Nexus5020(config)# interface vfc 1
Nexus5020(config-if)# bind interface ethernet 1/1
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 2
Nexus5020(config-if)# bind interface ethernet 1/2
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 3
Nexus5020(config-if)# bind interface ethernet 1/3
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 4
Nexus5020(config-if)# bind interface ethernet 1/4
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 5
Nexus5020(config-if)# bind interface ethernet 1/5
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 6
Nexus5020(config-if)# bind interface ethernet 1/6
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 10
Nexus5020(config-if)# bind interface ethernet 1/10
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 20
Nexus5020(config-if)# bind interface ethernet 1/20
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 21
Nexus5020(config-if)# bind interface ethernet 1/21
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 22
Nexus5020(config-if)# bind interface ethernet 1/22
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 23
Nexus5020(config-if)# bind interface ethernet 1/23
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 24
Nexus5020(config-if)# bind interface ethernet 1/24
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 25
Nexus5020(config-if)# bind interface ethernet 1/25
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 1-6
Nexus5020(config-if)# no shutdown
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 10
Nexus5020(config-if)# no shutdown
Nexus5020(config-if)# exit
Nexus5020(config)# interface vfc 20-25
Nexus5020(config-if)# no shutdown
Nexus5020(config-if)# exit
```

---

7. Enable the FC ports.  
In this example, FC ports 2/1 through 2/8 are enabled.

```
Nexus5010(config)# interface fc 2/1-8
Nexus5010(config-if)# no shutdown
Nexus5010(config-if)# exit
Nexus5010(config)# exit
```

8. Copy the running configuration to the startup configuration location.

```
Nexus5010# copy running-config startup-config
[#####] 100%
```

9. Copy the running configuration to a backup location.

```
Nexus5010# copy running-config ftp://10.10.20.1/backup.txt
```

10. Verify the configuration.

```
Nexus5010# show interface brief
```

```
Nexus5010# show running-config
```

## SAN scaling

When you expand a topology, avoid making changes that disrupt the original design goals. If data access requirements have changed, consider migrating to a topology that meets the current needs.

A high-availability SAN design accommodates a nondisruptive expansion. See [“SAN fabric topologies” \(page 23\)](#).

- 
- △ CAUTION:** HP highly recommends that you stop all I/O activity and back up all data before adding switches to the fabric.
- 

To expand a SAN:

- Use switches with a higher port count.
- Increase the number of switches.

To add switches to an existing fabric, ensure that the new configuration conforms to fabric rules. See the following fabric rules sections:

- [“Fabric rules” \(page 134\)](#), for H-series switches
  - [“Fibre Channel switch fabric rules” \(page 94\)](#), for B-series switches
  - [“Fibre Channel switch fabric rules” \(page 126\)](#), for C-series switches
  - [“Fibre Channel fabric rules” \(page 142\)](#), for M-series switches
- Add another fabric as a high-availability NSPOF solution.
  - Implement Fibre Channel routing. See [“Fibre Channel routing” \(page 43\)](#).
  - Deploy multiple independent SANs.
  - Migrate to a different SAN topology. See [“Topology migration” \(page 41\)](#).

When adding new switches, avoid fabric segmentation. For more information, see [“Fabric segmentation errors” \(page 391\)](#).

## Cascaded fabric expansion

Expand a cascaded fabric by connecting a new switch to an available port on an existing switch. If no ports are available, remove devices from an existing switch, connect the new switch to those ports, and then reconnect the devices to the new switch.

## Meshed fabric expansion

Expand a meshed fabric by connecting a new switch to available ports on an existing switch. If no ports are available, remove devices from an existing switch, connect the new switch to those ports, and then reconnect the devices to the new switch. Ensure that multiple paths (ISLs) connect the new switch to the meshed fabric.

## Ring fabric expansion

Expand a ring fabric by adding a switch to the ring.

Add new switches cascaded off the ring, up to the maximum number of switches supported in a single fabric. When expanding outside the ring, ensure that communicating devices are connected by no more than seven hops.

## Core-edge fabric expansion

Expand a core-edge SAN fabric by adding edge switches. Connect edge switches to available ports on the backbone switches.

If the current SAN contains only one core switch, add another. Connect edge switches to ports on the new core switch.

## SAN fabric merging

You can merge independent fabrics into a single, larger fabric. Merging enables you to:

- Provide more resources to independent SAN fabrics.
- Share the resources in two or more fabrics.
- Make information in one SAN available to servers in another SAN.
- Connect geographically dispersed fabrics into a single SAN.

During initial login, the discovery process determines whether the fabrics are compatible. If not, segmentation occurs and they continue to operate as separate fabrics. See [“Fabric segmentation errors” \(page 391\)](#).

When fabrics merge, the zone configuration databases are updated to include the zone configurations of each fabric. If a nonzoned fabric merges with a zoned fabric, all zoning information is proliferated to the nonzoned fabric switches. A zone configuration that is enabled at the time of the merge is also enabled on the nonzoned fabric switches. Devices in the nonzoned fabric are not accessible until you add them to the current configuration.

---

**NOTE:** When enabling a new configuration, HP strongly recommends that the fabric be quiesced. Zone membership should not be changed for devices that are actively performing I/O in a fabric. Once the new zoning is enabled, an RSCN is sent to all nodes that have been registered to receive the RSCN.

---

The following sections provide information to help you successfully merge SAN fabrics:

- [“Fabric segmentation errors” \(page 391\)](#)
- [“Switch configuration parameters” \(page 391\)](#)
- [“Independent fabric merge” \(page 391\)](#)
- [“High-availability redundant fabric merge” \(page 391\)](#)

## Fabric segmentation errors

The following errors can cause fabric segmentation:

- **Zone type mismatch**

The name of a zone object in one fabric is identical to the name of a different type of zone object in the other fabric. For example, an object name on fabric A must not be an alias or configuration name in fabric B; otherwise, the fabrics cannot merge.

- **Zone content mismatch**

The definition of a zone object in one fabric is different from its definition in the other fabric. If an alias, zone, or configuration name is the same on both fabric A and B, but the content or definition of the objects is different, the fabrics cannot merge.

- **Zone configuration mismatch**

Zoning is enabled for both fabrics, but the zone configurations are different. The fabrics cannot merge until the zone configuration is disabled for one of the fabrics.

- **Duplicate IDs**

Each switch in a fabric must have a unique domain ID; each switch in multiple fabrics of the SAN must also have a unique ID. For example, if fabric A has five switches with domain IDs 1 through 5, and fabric B has five switches with the same domain IDs, these two fabrics cannot merge until each switch in both fabrics has a unique domain ID.

For port-level zoning, changing the domain IDs can disrupt device access. Port-level zones are based on the domain ID and port number.

## Switch configuration parameters

Mismatched switch parameters can cause a SAN merger to fail. The configuration settings for all switches in the fabric must match, with the exception of the following parameters:

- Switch name
- IP address
- Domain ID

## Independent fabric merge

Merge fabrics by disabling the effective configuration on one fabric and then connecting both fabrics. After you connect the fabrics, devices in the second fabric are not accessible until you add them to the effective configuration.

---

**△ CAUTION:** When you disable the effective configuration, the fabric becomes accessible to all servers.

---

To merge two fabrics without disabling the effective configuration for entire fabrics, disable at least one switch in each fabric or use an additional switch. Use the disabled switch to merge the fabrics and create the new configuration.

## High-availability redundant fabric merge

With redundant fabrics, you can merge a fabric by taking it offline and redirecting I/O to the other fabric. Current I/O operations are not affected; however, during the merge, the hosts operate in degraded mode without redundant data path protection. With proper planning, you can minimize downtime. After completing and verifying the fabric merge, bring the first fabric online by restoring the I/O paths. After you restore the I/O paths on the first fabric, you can repeat the merge process for the second fabric.

### **Merging high-availability fabrics example**

In the following procedure, the SAN consists of fabric A and a redundant fabric B. Each of these fabrics is merged with a SAN consisting of fabrics C and D.

1. Identify and resolve any issues that can cause fabric segmentation.
2. Verify that each fabric provides a redundant path to all attached devices.
3. Verify that paths are open to each device that must remain online during the merge.
4. Select fabrics for merging, for example, fabric A with fabric C.
5. Close all active paths on the fabric selected for merging and prepare devices for downtime. For example, use multipathing software to redirect I/O by performing a failover to the alternate path.
6. Verify that the fabric selected for merging has no I/O activity.
7. Connect the selected fabrics (fabric A and fabric C).
8. Verify that the newly merged fabric contains all switches and that the zoning has merged correctly.
9. Restore I/O operations on the new fabric from the multipathing software console.
10. Verify that paths are open and restored for each device.
11. Ensure that all paths and I/O operations have been restored.

Repeat this procedure to merge fabric B with fabric D.



---

## 22 Support and other resources

### Contacting HP

#### 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 email notification of product enhancements, new driver versions, firmware updates, and other product resources.

#### Documentation feedback

HP welcomes your feedback.

To make comments and suggestions about product documentation, please send a message to [storagedocsfeedback@hp.com](mailto:storagedocsfeedback@hp.com). All submissions become the property of HP.

#### New and changed information in this edition

Changes from the August 29, 2011 version of this guide include:

- Chapter 4—Added CNA cable distances
- Chapter 5—Updated supported firmware and added support for HP SN6000B 16Gb Fibre Channel Switch
- Chapter 10—Clarified P2000 G3 FC support for HP-UX and Apple Mac OS X
- Chapters 11—Updated P2000 G3 and MSA2000 support tables

#### Related information

**Table 209 Related documentation**

Topic	Information source
Switches	For the latest information on B-series, C-series, H-series, and M-series switches and firmware versions, see the SAN Infrastructure website:

**Table 209 Related documentation** (continued)

Topic	Information source
	<a href="http://h18006.www1.hp.com/storage/saninfrastructure.html">http://h18006.www1.hp.com/storage/saninfrastructure.html</a>
Fabric interoperability	<ul style="list-style-type: none"> <li>• <i>HP StorageWorks Fabric Interoperability: Merging Fabrics Based on C-series and B-series Fibre Channel Switches Application Notes</i></li> <li>• <i>HP StorageWorks Fabric Interoperability: Merging Fabrics Based on C-series and M-series Fibre Channel Switches Application Notes</i></li> <li>• <i>HP StorageWorks Fabric Interoperability: Merging Fabrics Based on M-series and B-series Fibre Channel Switches Application Notes</i></li> </ul> <p>See these documents on the SAN Infrastructure website:  <a href="http://h18006.www1.hp.com/storage/saninfrastructure.html">http://h18006.www1.hp.com/storage/saninfrastructure.html</a></p>
EBS	<p>EBS documents, including the EBS compatibility matrix, are available on the HP StorageWorks Enterprise Backup Solution website:  <a href="http://www.hp.com/go/ebs">http://www.hp.com/go/ebs</a></p> <p>The EBS compatibility matrix is available on the Tape Compatibility and Tools website:  <a href="http://h18000.www1.hp.com/products/storageworks/tapecompatibility.html">http://h18000.www1.hp.com/products/storageworks/tapecompatibility.html</a></p>
Storage arrays	<p>For the latest information on storage arrays, see the Disk Storage Systems website:  <a href="http://h18006.www1.hp.com/storage/arrayssystems.html">http://h18006.www1.hp.com/storage/arrayssystems.html</a></p>

## Typographic conventions

**Table 210 Document conventions**

Convention	Uses
Blue text: <a href="#">Table 210 (page 394)</a>	Cross-reference links and email addresses
Blue, underlined text: <a href="http://www.hp.com">http://www.hp.com</a>	Website addresses
<b>Bold</b> text	<ul style="list-style-type: none"> <li>• Keys that are pressed</li> <li>• Text typed into a GUI element, such as a box</li> <li>• GUI elements that are clicked or selected, such as menu and list items, buttons, tabs, and check boxes</li> </ul>
<i>Italic</i> text	Text emphasis
Monospace text	<ul style="list-style-type: none"> <li>• File and directory names</li> <li>• System output</li> <li>• Code</li> <li>• Commands, their arguments, and argument values</li> </ul>
<i>Monospace, italic</i> text	<ul style="list-style-type: none"> <li>• Code variables</li> <li>• Command variables</li> </ul>
<b>Monospace, bold</b> text	Emphasized monospace text
.	Indication that the example continues.

 **WARNING!** Indicates that failure to follow directions could result in bodily harm or death.

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

---

ⓘ **IMPORTANT:** Provides clarifying information or specific instructions.

---

**NOTE:** Provides additional information.

---

💡 **TIP:** Provides helpful hints and shortcuts.

---

## Customer self repair

HP CSR programs allow you to repair your HP product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider. For North America, see the CSR website:

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

This product has no customer-replaceable components.

---

# Glossary

This glossary defines acronyms and terms used in this guide and is not a comprehensive glossary of computer terms.

## Symbols

**3DC** Three Data Center. A storage configuration that uses XP storage systems.

**3PO** Third-party option.

## A

**AAA** Authentication, Authorization, and Accounting. A security protocol.

**ABM** Array-based management.

**access authorization** A fabric security method that provides frame-level access control in hardware and verifies the SID-DID combination of each frame. Also known as *hard zoning*.

**ACL** Access control list. A data protection feature used to restrict access to data resources based on defined policies.

**ACU** Array Configuration Utility. A utility used to control LUN access.

**ADG** Advanced Data Guarding.

**AES** Advanced Encryption Standard. An encryption security technology.

**AG** Access Gateway. A software-enabled feature that allows B-series Blade switches to function as port aggregators.

**arbitrated loop** See FC-AL.

**ARL** Adaptive Rate Limiting.

**ASIC** Application-specific integrated circuit.

**ATM** Asynchronous Transfer Mode. A communications networking technology for LANs and WANs that carries information in fixed-size cells of 53 bytes (5 protocol bytes and 48 data bytes).

## B

**B-series** Fibre Channel switches manufactured for HP by Brocade Communications Systems, Inc.

**BB\_credits** Buffer-to-buffer credits.

**BCN** Backward Congestion Notification.

**BE** Back-side path. A path between the Data Path Module and the physical storage (for example, an HP EVA).  
Storage on the back-end array (SVSP).

**BFS** Boot from SAN.

## C

**C-series** Fibre Channel switches manufactured for HP by Cisco Systems, Inc.

**CA** Certificate authority.

**cascaded switch** The destination director attached to the control unit.

**CBC** Cipher-block chaining. A data encryption cipher mode.

**CEE** Converged Enhanced Ethernet.

**CFS** Cisco Fabric Services. A C-series port security feature.

**CHAP** Challenge-Handshake Authentication Protocol. An IP SAN security technology.

**CLI** Command line interface.

**CLIM** Cluster I/O module. A hardware component of the HP NonStop server.

<b>CN</b>	Converged Network.
<b>CNA</b>	Converged Network Adapter.
<b>controller pair</b>	Two connected controller modules that control a disk array.
<b>CWDM</b>	Coarse wavelength division multiplexing. The technique of placing more than one optical signal on a single optical cable simultaneously. <i>See also</i> WDM.

## D

<b>DA</b>	Directory agent. An iSCSI software component.
<b>DAS</b>	Direct-attached storage.
<b>DCB</b>	Data Center Bridging (Cisco).
<b>DCBX</b>	Data Center Bridging Exchange Protocol.
<b>DCC</b>	Device Connection Control. A data access policy that determines to which switch ports a Fibre Channel device can connect.
<b>DCE</b>	Data Center Ethernet.
<b>DCFM</b>	Data Center Fabric Manager. A B-series switch feature used to centralize fabric management through a host-based application.
<b>DES</b>	Data Encryption Standard. A block cipher designed for use in symmetric cryptography, which encrypts data in 64-bit blocks and uses a key length of 56 bits.
<b>DHCHAP</b>	Diffie-Hellman Challenge-Handshake Authentication Protocol. An IP SAN security technology.
<b>director fabric</b>	A SAN fabric that uses HP Director model switches.
<b>discovery authentication</b>	A security method where the fabric presents only a partial list of authorized devices. Also known as <i>soft zoning</i> .
<b>DLS</b>	Dynamic load sharing.
<b>DMP</b>	Dynamic Multipathing software.
<b>DMS</b>	Database Migration Services.
<b>DPM</b>	Data Path Module. A SAN-based device, separate from the core Fibre Channel switching infrastructure, that provides storage virtualization services across heterogeneous hosts, storage, and SAN fabrics.
<b>DSM</b>	Device Specific Module.
<b>DWDM</b>	Dense wavelength division multiplexing. The technique of placing several optical signals on a single optical cable simultaneously. <i>See also</i> WDM.

## E

<b>EBS</b>	Enterprise Backup Solution. HP SAN backup software.
<b>ECB</b>	Electronic code book. An encryption cipher mode.
<b>EE</b>	Encryption Engine. A feature of the B-series Encryption Switch.
<b>EFMD</b>	Exchange Fabric Membership Data. A protocol used in C-series fabric binding configurations.
<b>EFMS</b>	HP StorageWorks Enterprise Fabric Management Suite. Advanced switch management software for H-series switches.
<b>enterprise</b>	Any large organization where information technology is essential for continuing operations.
<b>EoR</b>	End of Row (switch).
<b>ESCON</b>	Enterprise Systems Connection.
<b>ESOL</b>	End Of Support Life.
<b>ESS</b>	Enterprise Storage System.

<b>EVA</b>	HP StorageWorks Enterprise Virtual Array. A high-performance, high-capacity, and high-availability storage solution for the high-end enterprise class marketplace. Each EVA storage system consists of a pair of HSV virtualizing storage controllers and the disk drives they manage.
<b>F</b>	
<b>fabric</b>	A network of one or more Fibre Channel switches that transmit data between any two N_Ports on any of the switches.
<b>failover</b>	An automatic method for transferring operations from a failed system to a secondary, identical system.
<b>Fat tree topology</b>	A SAN topology where at least 50% of edge ports are dedicated as ISLs.
<b>FC</b>	Fibre Channel. A comprehensive set of standards for concurrent communication among servers, storage systems, and peripheral devices.
<b>FC-AL</b>	Fibre Channel Arbitrated Loop. A serial, full-duplex, data transfer architecture for high-performance storage systems.
<b>FC-SONET</b>	Fibre Channel Synchronous Optical Network. A gateway that resides at the end of an inter-site link and encapsulates Fibre Channel frames into SONET packets before transmitting the frames over the network.
<b>FC-SP</b>	Fibre Channel Security Protocol. A security protocol that protects in-transit data.
<b>FC-to-ATM</b>	Fibre Channel to Asynchronous Transfer Mode.
<b>FC-to-IP</b>	Fibre Channel to Internet Protocol.
<b>FCC</b>	Fibre Channel Congestion Control. A feature that allows C-series switches to intelligently regulate traffic across ISLs and ensure that each initiator-target pair of devices has the required bandwidth for data transfer.
<b>FCDM</b>	Fibre Channel Disk Module.
<b>FCF</b>	Fibre Channel forwarder.
<b>FCFW</b>	Fibre Channel FastWrite.
<b>FCID</b>	Fibre Channel ID.
<b>FCIP</b>	Fibre Channel over Internet Protocol.
<b>FCoE</b>	Fibre Channel over Ethernet Protocol.
<b>FCS</b>	1. Fibre Channel switch. 2. Fabric Configuration Server, a data access policy that determines which switches can change fabric configurations.
<b>FCSA</b>	Fibre Channel ServerNet Adapter. The HP NonStop version of Fibre Channel HBAs, used to connect to XP disk arrays.
<b>FDISC</b>	Fabric discovery.
<b>FE</b>	FE front-side path. A path between the host (host bus adapter) and the DPM. Front-end host visible (SVSP).
<b>FEC</b>	Forward Error Correction.
<b>FF</b>	Full-featured.
<b>Fibre Channel</b>	A serial data transport infrastructure and protocol used to implement SANs. See also <a href="http://www.fibrechannel.org/">http://www.fibrechannel.org/</a> and <a href="http://www.t11.org/">http://www.t11.org/</a> .
<b>FIP</b>	FCoE Initialization Protocol.
<b>FLOGI</b>	Fabric login.
<b>FOS</b>	Fabric Operating System.
<b>FPMA</b>	Fabric Provided MAC Addressing.
<b>FRU</b>	Field-replaceable unit.

## G

<b>Gb/s</b>	Gigabits per second.
<b>GbE</b>	Gigabit Ethernet. An Ethernet standard for transmitting data at 1 Gb/s.

<b>GBIC</b>	Gigabit interface converter. A hardware module that connects fiber optic cables to a device and converts electrical signals to optical signals.
<b>GLM</b>	Gigabit link module. A 1 Gb/s fiber optic transceiver.
<b>GPS</b>	General-purpose server.
<b>H</b>	
<b>H-series</b>	Fibre Channel switches manufactured for HP by QLogic.
<b>HA</b>	High availability. The relative ability of a system to operate continuously, regardless of the type of failure.
<b>HACMP/ES</b>	High Availability Cluster Multiprocessing/Enhanced Scalability.
<b>HAFM</b>	High Availability Fabric Manager. A software tool used to configure and manage fabrics.
<b>HBA</b>	Host bus adapter. A hardware device that connects the host server to the fabric.
<b>heterogeneous</b>	A mixed environment that incorporates different operating systems, protocols, architectures, and equipment from different vendors or product families.
<b>high availability</b>	See HA.
<b>HLDM</b>	Hitachi Dynamic Link Manager.
<b>hop</b>	An interswitch link between a pair of Fibre Channel switches.
<b>host bus adapter</b>	See HBA.
<b>HotCAT</b>	Hot Code Activation Technology. An M-series switch feature that provides nondisruptive code activation.
<b>HP Network Advisor</b>	Host-based application that centralizes fabric management.
<b>HP P6000 Continuous Access</b>	An HP storage product consisting of two or more EVA storage arrays performing disk-to-disk replication, and a management user interface that facilitates configuring, monitoring, and maintaining replication on the storage systems.
<b>HP P9000 (XP) Continuous Access</b>	An HP storage product consisting of two or more XP disk arrays performing disk-to-disk replication, and a management user interface that facilitates configuring, monitoring, and maintaining replication of the storage systems.
<b>HTTPS</b>	Hypertext Transfer Protocol over Secure Socket Layer.
<b>I</b>	
<b>ICL</b>	Inter-chassis link. A feature that provides dedicated high-bandwidth links between two DC SAN Backbone Director chassis.
<b>iFCP</b>	Internet Fibre Channel Protocol.
<b>IFL</b>	Inter-fabric link. A pair of switches that provide redundant paths between fabrics.
<b>IFR</b>	Inter-fabric routing. A feature used to enable device sharing across Virtual Fabrics or VSANs.
<b>IFZ</b>	Inter-fabric zone. A component in inter-fabric routing that contains exactly three WWN members: the local device, the remote device, and the TR_Port attached to the remote fabric.
<b>iLO</b>	Integrated Lights Out. An HP ProLiant Storage Server connectivity feature.
<b>in-band communication</b>	Communication through the same communications channel as the operational data.
<b>INCITS</b>	International Committee for Information Technology Standards.
<b>interswitch link</b>	See ISL.
<b>IOA</b>	IO Accelerator.
<b>IOAME</b>	I/O adapter module enclosure.
<b>IOS</b>	Internetwork Operating System (Cisco).
<b>IPS</b>	C-series Internet Protocol Storage Services Module.
<b>IPsec</b>	Internet Protocol Security. A method of data encryption.

<b>iSCSI</b>	Internet Small Computer System Interface. A standard protocol that uses SCSI commands to transfer data over IP networks.
<b>ISID</b>	Initiator session ID.
<b>ISL</b>	Interswitch link. A connection from an E-port on one switch to an E-port on another switch.
<b>iSNS</b>	Internet Storage Name Service. An iSCSI client-server discovery protocol.
<b>IT</b>	Information technology.
<b>IVR</b>	Inter-VSAN routing. A C-series topology that provides selective Fibre Channel routing connectivity between devices in different VSANs.

## L

<b>LAG</b>	Link aggregation group.
<b>LC</b>	Lucent connector. An industry-standard connector for fiber optic cable connections.
<b>LDEV</b>	Logical device.
<b>LFF</b>	Large form factor. Used to describe a type of drive cage.
<b>LLDP</b>	Link Layer Discovery Protocol.
<b>LR</b>	Long Reach.
<b>LSAN</b>	Logical SAN. An interconnected network of independent fabrics for B-series switches.
<b>LTO</b>	Linear Tape-Open.
<b>LUN</b>	Logical unit number. An identification scheme for storage disks.
<b>LUSE</b>	Logical unit size expansion.

## M

<b>M-series</b>	Fibre Channel switches manufactured for HP by McDATA Corporation.
<b>MAN</b>	Metropolitan area network. A communications network that covers a geographic area such as a town, city, or suburb.
<b>MCU</b>	Master control unit.
<b>MIB</b>	Management Information Base (B-series).
<b>MME</b>	Multimedia encryption.
<b>MP Router</b>	Multi-protocol Router.
<b>MPIO</b>	Multipath I/O (Microsoft software).
<b>MSA</b>	Modular Smart Array. An HP storage system.
<b>MSCS</b>	Microsoft Cluster Server software.
<b>MSM</b>	C-series Multiservice Module.
<b>MTU</b>	Maximum transmission unit.

## N

<b>NAT</b>	Network address translation.
<b>NDCLA</b>	Nondisruptive code load and activation. A high-availability switch feature.
<b>NIC</b>	Network interface card.
<b>NPIV</b>	N_Port ID Virtualization. An industry-standard protocol.
<b>NPV</b>	N_Port Virtualization.
<b>NSPOF</b>	No single point of failure. A configuration where failure of a single component does not cause failure of the entire system.

## O

<b>OADM</b>	Optical add-drop multiplexer.
-------------	-------------------------------



<b>OC</b>	Optical Carrier. A range of digital signals used on a SONET network.
<b>OCP</b>	Operator control panel. A panel on the EVA controller that you use to manage the EVA.
<b>OS</b>	Operating system.
<b>out-of-band communication</b>	Communication through a different communications channel than that used by operational data.
<b>oversubscription</b>	When transmissions from one or more devices exceed the capacity of an ISL or port.
<b>P</b>	
<b>PAP</b>	Password Authentication Protocol. A security protocol.
<b>PCM</b>	Power and control modules.
<b>PDC</b>	Processor-Dependent Code.
<b>PDU</b>	Protocol data unit.
<b>PFC</b>	Priority flow control.
<b>PIC</b>	Plug-in card. A hardware component of the HP NonStop server.
<b>PKI</b>	Public key infrastructure. A data encryption technology.
<b>platform</b>	A combination of supported hardware components and operating system on a server from a specific vendor.
<b>pWWN</b>	Port World Wide Name. Synonymous with WWPN.
<b>Q</b>	
<b>QoS</b>	Quality of service. A traffic management feature.
<b>R</b>	
<b>RADIUS</b>	Remote Authentication Dial-In User Service. A client-server protocol that protects networks against unauthorized access.
<b>RBAC</b>	Role-based access control. A Fabric OS feature used to determine which commands are supported for each user.
<b>RCS</b>	Remote copy set.
<b>RCU</b>	Remote control unit.
<b>RHEL</b>	Red Hat Enterprise Linux.
<b>RPO</b>	Recovery point objective. A disaster recovery concept that defines what an organization considers as an acceptable loss in terms of time.
<b>RSPAN</b>	Remote Switched Port Analyzer.
<b>RTO</b>	Recovery time objective. A disaster recovery determination of when the recovery site is to be used.
<b>RTT</b>	Round-trip time.
<b>RU</b>	Rack unit.
<b>RVU</b>	Release version updates.
<b>S</b>	
<b>SA</b>	Service Agent. An iSCSI software component.
<b>SAN</b>	Storage area network. An intelligent infrastructure that connects heterogeneous servers with shared, heterogeneous storage systems.
<b>SAN-OS</b>	SAN Operating System (Cisco).
<b>SAS</b>	Serial Attached SCSI. A bus technology used to transfer storage data.
<b>SATA</b>	Serial Advanced Technology Attachment. A bus technology for storage devices.
<b>SC</b>	Subscriber connector.

<b>SCC</b>	Switch Connection Control. A data access policy that determines to which switches a switch can join.
<b>SCP</b>	Secure Copy. A B-series feature used to securely transfer files between systems.
<b>SFF</b>	Small form factor. Used to describe a type of drive cage.
<b>SFP</b>	Small form-factor pluggable transceiver.
<b>SFP+</b>	Small form-factor pluggable, plus transceiver.
<b>SFTP</b>	Secured File Transfer Protocol. A C-series security feature.
<b>SID</b>	Source ID.
<b>Skinny tree topology</b>	A SAN topology where less than 50% of the edge ports are dedicated as ISLs.
<b>SKU</b>	Stock-keeping unit.
<b>SLES</b>	SUSE Linux Enterprise Server.
<b>SMA</b>	HP StorageWorks Storage Management Appliance.
<b>SMB</b>	Small and medium business. Any organization that uses onsite computer systems.
<b>SME</b>	Storage Media Encryption (C-series). A standards-based encryption solution for heterogeneous and virtual tape libraries.
<b>SNIA</b>	Storage Networking Industry Association.
<b>SNS</b>	Simple Name Service. A fabric service that provides a mapping between device names and their addresses in a fabric.
<b>SONET</b>	Synchronous Optical Network. A network used to create an extended SAN by connecting local SANs.
<b>SP</b>	Service pack. A collection of software updates.
<b>SPAN</b>	Switched Port Analyzer. A troubleshooting tool.
<b>SPARC</b>	Scalable Processor Architecture. An instruction-set architecture typically used with Sun systems.
<b>SPOCK</b>	Single Point of Connectivity Knowledge. An HP website as the primary portal to obtain detailed information about supported HP storage product configurations. An HP Passport account is required for access.
<b>SPOF</b>	Single point of failure.
<b>SR</b>	Short Reach.
<b>SRM</b>	Storage resource management.
<b>SSCM</b>	HP StorageWorks Simple SAN Connection Manager. Comprehensive management software for H-series switches.
<b>SSH</b>	Secure Shell. A security encryption protocol.
<b>SSIF</b>	Storage Security Industry Forum.
<b>SSL</b>	Secure Socket Layer. A security protocol that provides access to a fabric through web-based management tools.
<b>SSM</b>	Storage Services Module (C-series).
<b>SSP</b>	Selective Storage Presentation. A feature used to restrict access to a Fibre Channel LUN.
<b>ST</b>	Straight-tip connector.
<b>SVSP</b>	SAN Virtualization Services Platform.
<b>SWCC</b>	StorageWorks Command Console.
<b>T</b>	
<b>TACACS+</b>	Terminal Access Controller Access Control System Plus. A client-server protocol.
<b>TDM</b>	Time-division multiplexing.
<b>TOE</b>	TCP Offload Engine.
<b>topology</b>	The physical structure of interconnected components that form a network.

<b>ToR</b>	Top of Rack (switch).
<b>TR</b>	Transparent router. A routed fabric feature that provides inter-fabric routing, allowing controlled and limited access between devices on an H-series switch (local) fabric and devices on a remote fabric of other vendor switches.
<b>TSID</b>	Target session ID.
<b>U</b>	
<b>UA</b>	User Agent. An iSCSI software component.
<b>unified fabric</b>	IEEE DCB enabled network (Cisco).
<b>USD</b>	UltraNet Storage Director (Brocade-McDATA firmware).
<b>V</b>	
<b>VAR</b>	Value-added reseller.
<b>VC-Enet</b>	Virtual Connect Ethernet. A hardware component of HP Virtual Connect.
<b>VC-FC</b>	Virtual Connect Fibre Channel. An optional Fibre Channel interconnect module for HP BladeSystem c-Class servers.
<b>VCS</b>	Virtual Controller Software. The software that runs in the HSV controller-based Enterprise Virtual Array GL model storage systems.
<b>VFC</b>	Virtual Fibre Channel.
<b>VIO</b>	Versatile I/O.
<b>VLAN</b>	Virtual logical area network.
<b>VLS</b>	Virtual Library System.
<b>VM</b>	Virtual machine.
<b>VOQ</b>	Virtual output queuing. A C-series feature used for nonblocking architecture.
<b>VOS</b>	Virtual operating system.
<b>VPN</b>	Virtual private network.
<b>VSAN</b>	Virtual storage area network. A logical SAN partition that can be configured and managed independently.
<b>VSM</b>	Virtualization Services Manager. Performs functions such as management and configuration of the SVSP as well as data copy features (migration, clone and async replication).
<b>VSS</b>	Volume Shadow Copy Service. Microsoft software used to make snapshots.
<b>W</b>	
<b>WDM</b>	Wavelength division multiplexing. The technique of placing multiple optical signals on a single optical cable simultaneously. <i>See also</i> CWDM and DWDM.
<b>WWN</b>	World Wide Name. A unique identifier assigned to a Fibre Channel device.
<b>WWPN</b>	World Wide Port Name. Synonymous with pWWN.
<b>X</b>	
<b>XCS</b>	Xcelerated Controller Software. The software that runs in the HSV controller-based Enterprise Virtual Array XL model storage systems.
<b>XISL</b>	A special Brocade ISL that can carry traffic for multiple Logical Fabrics while maintaining traffic separation for each Logical Fabric.
<b>Z</b>	
<b>zone</b>	A collection of devices or user ports that communicate with each other through a fabric. Any two devices or user ports that are not members of at least one common zone <i>cannot</i> communicate through the fabric.

# Index

## Symbols

3PAR *see* HP 3PAR

## A

access authorization, hard zoning enforcement, 376

Access Gateway mode

B-series switches, 88

c-Class BladeSystem, 163

configuration example, 163

considerations, 164

encryption, 90

failover and failback policies, 164

guidelines, 165

active protocol handling, 254

active signal amplification, 254

Adaptive Networking with QoS

B-series switches, 89

adaptive rate limiting

B-series switches, 89

addresses, importing and exporting, 44

advanced encryption standard, 361

Advanced Performance Monitor

B-series switches, 88

Advanced WebTools

B-series switches, 88

advanced zoning

B-series switches, 88

AES, 361

applications, zoning method, 377

array management

SVSP, 234

ATM extension Fibre Channel distance rules, 153

## B

B-series

1606 Extension SAN Switch, 280

400 MP Router, 283

iSCSI Director Blade, 341

scalability rules, 343

MP Router, 285

iSCSI support, 344

B-series features

Administrative Domains, 36

IFR, 36

Virtual Fabrics, 35

Virtual Fabrics with IFR, 36

B-series switches

database size, 98

Director, 84

domain ID, 98

edge, 84

extended fabric settings, 256

fabric groups, 47

Fabric Manager versions, 84, 85

fabric rules, 95

features, 88

firmware update, 84

firmware versions, 84, 85

ISL maximums, 98

LSANs, 47

maximum topology, 33

merging, 47

Meta SAN, 35

models, 84

overview, 83

routing, 48

SAN extension, 250

usage, 93

zoning enforcement, 112

zoning guidelines, 113

backup

SAN, 247

SAN based, 21

bandwidth

considerations, 261

determining requirements, 262

benefits

cascaded fabric, 26

core-edge fabric, 32

meshed fabric, 27

Meta SAN, 36

ring fabric, 29

SAN fabric, 24

single-switch fabric, 25

VSAN, 37

best practices

SAN, 372

storage security, 363

blade servers, supported switches, 161

booting

EVA storage systems, 221

HP 3PAR storage systems, 246

P9000/XP/VA, 230

buffer-to-buffer credits, extended, 251

Business Copy

SVSP support, 235

## C

c-Class BladeSystem

Brocade 4 Gb SAN Switch, 87, 91, 113

Brocade Access Gateway, 163

C-series, 118

FlexAttach, 168

supported blade servers, 161

Virtual Connect, 161

C-series features

IVR, 37

VSANS with IVR, 37

C-series Fibre Channel and FCoE switches

overview, 115

C-series switches

- capabilities, 115
- dividing, 48
- extended fabric settings, 257
- fabric rules, 127
- features, 123
- firmware versions, 118
- iSCSI, 345
- ISL maximums, 128
- maximum topology, 34
- MDS 9216i, IPS-4 IP Storage Services modules, IPS-8, and 14/2 services modules, 287
- models, 117
- routing, 50
- SAN extension, 250
- switching module support matrix, 120, 121
- usage, 125
- VSAN, 47
- VSAN configuration, 129
- zoning enforcement, 129
- cables
  - dressing, 375
  - high-availability options for P6000, 222
  - labeling, 374
  - record keeping, 374
  - tie wraps cautions, 375
- cascaded fabric
  - benefits, 26
  - expansion, 389
  - migrating from, 42
  - overview, 25
- case sensitivity, fabric identifiers, 378
- certified
  - FCIP vendor products, 311
- CHAP, 359
- Citrix Xen
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 126
- coarse wave division multiplexing *see* CWDM
- Compaq switch models
  - default settings (B-series), 94
- components
  - interconnect rules, 146
  - SAN, 18
- confidentiality
  - data encryption, 362
  - data separation, 362
  - traffic separation, 362
- configuration rules
  - B-series Fibre Channel switches, 94
  - C-series Fibre Channel switches, 126
  - ESX, 199
  - EVA, 214
  - HP 3PAR, 243
  - HP OpenVMS, 185
  - HP Tru64 UNIX, 187
  - HP-UX, 183
  - IBM AIX, 190
  - Linux, 192
  - M-series switches, 142
  - Mac OS X, 188
  - Microsoft Windows, 194
  - MSA1000, 210
  - MSA1500, 208
  - MSA1510i, 318
  - P6000, 214
  - P9000/XP storage systems, 228
  - routing, 52
  - storage management server, 222
  - Sun Solaris, 197
  - VA storage systems, 228
  - Xen, 200
- configuration, SAN best practices, 374
- configurations
  - data migration, 261
  - mirrored FCIP SAN, 261
  - shared link, 260, 264
  - unsupported, 57
- congestion, SAN, 20, 156
- connections, record keeping, 374
- connectivity
  - increasing, 41
  - ports, 21
- Continuous Access
  - SVSP support, 235
- conventions
  - document, 394
  - text symbols, 394
- converged network, 24, 61
- Converged network switches
  - C-series, 76
  - features, 78
  - high-availability feature comparison, 78
  - C-series expansion modules, 77
- core switch, 19
- core-edge fabric
  - benefits, 32
  - expansion, 390
  - H-series switches, 131
  - numeric representation, 31
  - overview, 29
  - recommended ISL ratios, 31
  - routing connection, 52
  - tree topologies, 30
  - types, 30
  - uses, 29
- custom SAN design, HP recommendations, 20
- customer self repair, 395
- CWDM, 254

## D

- data access types, 32
- data access, in SAN fabric topologies, 32
- data availability
  - design considerations, 41
  - factors affecting, 38
  - level 1, 38
  - level 2, 39
  - level 3, 39

- level 4, 40
- NSPOF, 40
- SAN, 38
- security, 362
- Data Center Fabric Manager *see* DCFM
- data encryption standard, 360
- data integrity, 362
- data migration, 261
  - HP 3PAR, 243
  - HP P9000, 231
- data protection
  - HP SVSP, 234
- data provisioning
  - HP SVSP, 234
- data replication
  - HP SVSP, 234
- database size, B-series switches, 98
- DC Dir Switch MP Extension Blade, 280
- DCFM
  - B-series switches, 88
- DES, 360
- design considerations
  - congestion, 20
  - data availability, 41
  - geographic layout, 20
  - interoperability, 20
  - migration, 20, 41
  - oversubscription, 20
  - performance workload, 20
  - scalability, 20
- design rules
  - B-series Fibre Channel switches, 94
  - C-series Fibre Channel switches, 126
  - M-series switches, 142
- designs, SAN, 20
- device port interfaces, 146
- director switch, 19
  - B-series, 84
- disaster tolerance, SAN, 21
- discovery authentication, soft zoning enforcement, 376
- discovery mechanisms, 313
- DL380 G4, 355
- DL580 G2, 355
- DL585, 355
- document
  - conventions, 394
- documentation
  - providing feedback, 393
  - related, 393
- domain ID
  - B-series Fibre Channel switches, 98
  - configuration settings, 391
  - M-series switches, 143
  - zone member identifier, 378

## E

- EBS
  - SAN infrastructure, 247
  - zoning, 377

- edge switch, 19
  - B-series, 84
  - B-series port usage, 94
- EFMS
  - features, 133
- encryption
  - security technologies, 360
- Enterprise Backup Solution *see* EBS
- Ethernet
  - Fibre Channel over, 61
- EVA storage systems
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 126
  - HP P6000 Continuous Access support, 289, 291
  - HP P6000 Continuous Accesssupport, 293
  - iSCSI Connectivity Option, 333
  - M-series fabric rules, 143
  - maximum configurations, 216
  - rules, 213
  - tape storage, 221
  - with mpx110 rules, 275
  - with MPX200 rules, 279
  - zoning, 221
- Expansion modules
  - for C-series CN switches, 77
- extended fabric
  - B-series switches, 88
  - settings, 256
- extended links, buffer-to-buffer credits, 251

## F

- fabric
  - B-series Fibre Channel switch rules, 94
  - C-series rules, 126
  - definition, 19
  - design rules, 146
  - groups, 47
  - H-series switches, 134
  - heterogeneous interoperable, 154
  - interoperable, 155
  - isolated, 44
  - M-series rules, 142
  - merging redundant, 391
  - performance, 156
  - redundancy, 51
  - routing, 51
  - size limits, 45
  - topology types, 23
- Fabric Manager
  - B-series switches, 88
- fabric partitioning
  - B-series, 48
  - C-series, 48
- fabric rules
  - FCoE switch, 78
- fabric services
  - limits, 45
  - separate, 44
- fabric watch

- B-series switches, 88
- fat tree, 30
- FC-SONET gateway, 264
- FC-SP, 360
- FC-to-ATM, 265
- FC4-16IP *see* B-series iSCSI Director Blade
- FCIP
  - B-series switches, 90
  - managing bandwidth, 261
  - protocol, 259
  - routing, 57
  - third-party QoS products, 264
- FCIP extension, Fibre Channel distance rules, 153
- FCoE
  - described, 61
  - end-to-end solution, 61
  - fabric-edge solution, 61
  - HP products, 61
  - HP solutions, 61
  - technology, 62
- FCoE SAN fabrics, 24
- FCoE switches
  - C-series, 76
    - high-availability feature comparison, 78
  - C-series expansion modules, 77
  - C-series fabric rules, 78
  - C-series features, 78
- FCoE technology, 24, 38
- fiber optic cables
  - interconnect rules, 146
  - loss budgets, 148
- fiber optic transceiver, 251
- Fibre Channel
  - distance rules, 151
  - iSCSI comparison, 312
  - long-distance technologies, 251
  - over Internet protocol, 259
  - routing, 43
  - routing distance rules, 154
  - routing implementations, 46
  - routing techniques, 46
  - routing, B-series switches, 90
  - technology overview, 18
- Fibre Channel over Ethernet *see* FCoE
- Fibre Channel SAN security technologies, 360
- Fibre Channel switch models
  - C-series, 117
- Fibre Channel switches
  - function, 19
  - interface, 146
- firmware updates
  - B-series Fibre Channel switches, 84
  - M-series switches, 143
- firmware versions
  - B-series Fibre Channel switches, 84, 85
  - C-series Fibre Channel switches, 118
  - Fabric Manager, 84, 85
  - H-series switches, 132
  - M-series switches, 140

- FlexAttach, 168

## G

- GBIC, 251
- geographic layout, SAN fabric topologies, 20

## H

- H-series
  - maximum topology, 34
  - port topology, 34
- H-series switches
  - extended fabric settings, 257
  - fabric routing, 48
  - fabric rules, 134, 135
  - features, 132
  - firmware update, 132
  - firmware version, 132
  - high-availability features, 134
  - ISL maximums, 136
  - models, 132
  - overview, 131
  - TR feature, 56
  - usage, 134
  - with TR, 37
  - zoning enforcement, 138
  - zoning limits, 138
- HA Fabric Manager (HAFM)
  - firmware version, 85
- hard zoning enforcement, 376
- HBA NPIV, 170
  - connectivity guidelines, 170
  - considerations, 170
- HBA's
  - common, 201
  - dual channel configuration, 225
  - labeling, 374
  - zoning method, 377
- help
  - obtaining, 393
- heterogeneous SAN
  - fabric design configuration rules, 146
  - HP 3PAR, 241
  - integration with HP P6000 Continuous Access, 220
  - MSA1000, 210
  - MSA20000fc support, 160
  - platforms and operating systems, 21
  - server configuration rules, 159
- high availability
  - and merging redundant fabrics, 391
  - C-series VSAN, 129
  - data availability types, 40
  - EVA, 222
  - MP Router configurations, 55
  - routed SAN, 51
- high-availability features
  - B-series switches, 91
- hops
  - count for 400 MP Router and MP Router Blade, 106
  - limits per fabric, 26

- M-series maximum, [143](#)
  - minimizing counts, [21](#)
  - MP Router, [111](#)
  - host-based volume shadowing, OpenVMS, [186](#), [310](#)
  - host-based zoning, [376](#)
  - HP
    - B-series Fibre Channel switches, [85](#)
    - BladeSystem c3000 enclosure, [169](#)
    - C-series Fibre Channel switches, [124](#)
    - M-series switches, [140](#)
    - SAN Virtualization Services Platform, [234](#)
    - standard SAN designs, [20](#)
    - technical support, [393](#)
  - HP 3PAR
    - configuration rules, [243](#)
    - data migration, [243](#)
    - heterogeneous SAN support, [241](#)
    - storage management server integration, [246](#)
    - storage system rules, [241](#)
    - zoning, [246](#)
  - HP IP Distance Gateway, [266](#)
  - HP MPX200 Multifunction Router *see* MPX200
  - HP OpenVMS
    - B-series Fibre Channel switch fabric rules, [95](#)
    - C-series Fibre Channel switch fabric rules, [126](#)
    - configuration rules, [185](#)
    - host-based volume shadowing, [186](#), [310](#)
    - M-series fabric rules, [143](#)
    - multipathing coexistence support, [186](#)
    - SAN rules, [185](#)
    - supported storage systems, [185](#)
  - HP P6000 Command View
    - MPX200 support, [329](#)
  - HP P6000 Continuous Access
    - SAN extension best practices, [289](#)
    - SAN integration, [220](#)
    - support, [289](#), [291](#), [293](#)
  - HP P9000
    - data migration, [231](#)
  - HP P9000 (XP) Continuous Access, [296](#)
    - asynchronous replication rules, [298](#)
    - configuration support, [300](#)
    - direct storage-to-storage distances, [301](#)
    - ESCON director and repeater distances, [303](#)
    - ESCON IP gateway support, [310](#)
    - FCIP gateway support, [309](#)
    - Fibre Channel distances, [302](#)
    - iFCP gateway support, [310](#)
    - journal replication rules, [299](#)
    - synchronous replication rules, [298](#)
    - WDM distances, [303](#)
  - HP ProLiant Storage Server, iSCSI feature pack, [349](#)
  - HP Secure Advantage, [361](#)
  - HP SVSP, [234](#)
  - HP SVSP Business Copy, [235](#)
  - HP SVSP Continuous Access, [235](#)
  - HP SVSP thin provisioning, [235](#)
  - HP SVSP Volume Manager, [235](#)
  - HP Tru64 UNIX
    - B-series Fibre Channel switch fabric rules, [95](#)
    - C-series Fibre Channel switch fabric rules, [126](#)
    - configuration rules, [187](#)
    - M-series fabric rules, [143](#)
    - multipathing coexistence support, [188](#)
    - SAN rules, [187](#)
    - supported storage systems, [187](#)
  - HP Virtualization Services Manager, [235](#)
  - HP Virtualization Services Manager Command Line Interface, [235](#)
  - HP-UX
    - B-series Fibre Channel switch fabric rules, [95](#)
    - C-series Fibre Channel switch fabric rules, [126](#)
    - configuration rules, [183](#)
    - M-series fabric rules, [143](#)
    - multipathing coexistence support, [184](#)
    - SAN rules, [183](#)
    - supported storage systems, [183](#)
- ## I
- IBM AIX
    - B-series Fibre Channel switch fabric rules, [95](#)
    - C-series Fibre Channel switch fabric rules, [126](#)
    - configuration rules, [190](#)
    - M-series fabric rules, [143](#)
    - multipathing coexistence support, [191](#)
    - SAN rules, [190](#)
    - supported storage systems, [190](#)
  - ICL
    - DC/DC04 SAN Directors, [89](#)
  - IDs, duplicate, [391](#)
  - integrated iSCSI connectivity, [324](#)
  - inter-fabric zone, [50](#)
  - Internet Storage Name Service, [314](#)
  - interoperability, design considerations, [20](#)
  - IP address, configuration settings, [391](#)
  - IP SAN security technologies, [359](#)
  - IPsec, [360](#)
  - iSCSI
    - bridge, [323](#)
    - bridging to Fibre Channel, [323](#)
    - clustering, [351](#)
    - comparison with Fibre Channel, [312](#)
    - concepts, [312](#)
    - configuration rules, [348](#)
    - connectivity, [324](#)
    - discovery, [314](#)
    - distance rules for bridging Fibre Channel, [154](#)
    - enabled storage, [323](#)
    - initiators, [313](#), [315](#)
    - license upgrades, [350](#)
    - Microsoft Exchange solution, [351](#)
    - names, [313](#)
    - native products, [316](#)
    - overview, [312](#)
    - routing, [323](#)
    - security, [315](#)
    - sessions and logins, [314](#)
    - storage, [312](#)



- target, 313
- iSCSI Connectivity Option
  - EVA and EVA4400, 333
  - hardware support, 335
  - iSCSI boot, 339
  - rules, 340
  - software support, 336
  - supported maximums, 340
  - supported NICs, 339
  - supported operating systems, 339
- iSCSI feature pack
  - HP ProLiant Storage Server, 349
- ISL
  - B-series Fibre Channel switch maximum, 98
  - C-series maximum, 128
  - H-series switch maximum, 136
  - M-series maximum, 144
  - ratios for core-edge fabric, 31
  - SAN extension, 249
- iSNS, 314
- IVR
  - basic configuration, 47
  - overview, 37
- K**
- key management, 361
- L**
- labels
  - cable, 374
  - HBA, 374
- latency
  - SAN, 156
- level 4 high availability, EVA, 222
- Linux
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 126
  - configuration rules, 192
  - H-series Fibre Channel Switch fabric rules, 135
  - M-series fabric rules, 143
  - Red Hat supported storage systems, 192
  - SAN rules, 192
  - SUSE supported storage systems, 192
- login authentication, soft-plus zoning enforcement, 376
- LSAN, 35, 47, 48
- LUN masking, storage-based, 375
- M**
- M-series switches
  - Director, 140
  - domain ID, 143
  - EOSL, 140
  - extended fabric settings, 258
  - fabric rules, 143
  - features, 142
  - firmware update, 143
  - firmware versions, 140
  - hops, 143
  - ISL maximums, 144
  - maximum quantity, 143
  - maximum topology, 34
  - models, 140
  - operating systems supported, 142
  - overview, 140
  - port quantities, 140
  - SAN extension, 250
  - storage products, 142
  - usage, 142
  - user port maximum quantity, 143
  - zoning enforcement, 145
- Mac OS X
  - B-series Fibre Channel switch fabric rules, 95, 126
  - configuration rules, 188
  - supported storage systems, 188
- maximum configurations
  - EVA storage systems, 216
  - MSA1000, 211
  - MSA1500, 208
- maximum sustained I/O load, 260, 264
- maximums
  - C-series Fibre Channel switch fabric, 127
  - for SAN fabric topologies, 33
  - iSCSI Connectivity Option, 340
  - M-series fabric, 143
- McDATA
  - models, 140
- McDATA switches
  - for p-Class BladeSystems, 140
- mergers
  - connecting independent fabrics, 391
  - fabric, 390
  - high-availability redundant fabrics, 391
- meshed fabric
  - benefits, 27
  - expansion, 390
  - migrating from, 42
  - overview, 26
- Meta SAN, 47
  - B-series switches, 35
  - benefits, 36
  - fabric topologies, 36
  - overview, 35
  - scalability rules, 110
- Microsoft Exchange Server 2003, iSCSI, 351
- Microsoft Windows
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 126
  - configuration rules, 194
  - H-series Fibre Channel Switch fabric rules, 135
  - M-series fabric rules, 143
  - multipathing coexistence support, 195
  - SAN rules, 193
  - supported storage systems, 193
- migration
  - cascaded fabric, 42
  - cascaded to core-edge fabric, 42
  - cascaded to meshed fabric, 42
  - cascaded to ring fabric, 42

- meshed fabric, [42](#)
- meshed to core-edge fabric, [42](#)
- meshed to ring fabric, [42](#)
- nondisruptive, [41](#)
- ring fabric, [42](#)
- ring to core-edge fabric, [42](#)
- ring to meshed fabric, [42](#)
- topology, [41](#)
- mirrored FCIP SAN, [261](#)
- mismatches
  - switch parameters, [391](#)
  - zone configuration, [391](#)
- model names
  - B-series Fibre Channel switches, [84](#)
  - C-series Fibre Channel switches, [117](#)
  - FCoE switches, [76](#)
  - H-series switches, [132](#)
  - M-series switches, [140](#)
- model numbers
  - B-series Fibre Channel switches, [84](#)
  - C-series switches, [117](#)
  - H-series switches, [132](#)
- modified HP standard SAN designs, [20](#)
- Modular Smart Array *see* MSA
- MP Router
  - backbone fabric, [111](#)
  - basic configuration, [47](#)
  - fabric rules, [107](#)
  - firmware version, [84](#)
  - high-availability configuration, [55](#)
  - hop count, [111](#)
  - IP subnets, [287](#)
  - maximum hop count, [111](#)
  - overview, [83](#)
  - scalability, [110](#)
  - scalability rules, [110](#)
  - unsupported configurations, [57](#)
  - use case configurations, [57](#)
- MP Router Blade, [283](#)
- mpx110, [266](#)
- mpx110 gateway
  - configuration rules, [274](#)
- MPX200, [324](#)
  - configuration options, [324](#)
  - configuration rules for FCIP, [279](#)
  - hardware support, [328](#)
  - iSCSI rules, [328](#)
  - management rules, [329](#)
  - multipath software, [330](#)
  - switch support, [330](#)
  - using iSCSI, [324](#)
  - with FCIP, [276](#)
  - with HP 3PAR, [243](#)
  - with HP P6000 Command View, [329](#)
  - with P6000/EVA, [329](#)
- MSA storage systems
  - C-series Fibre Channel switch fabric rules, [126](#)
  - M-series fabric rules, [143](#)
- MSA1000
  - configuration rules, [208, 210, 211](#)
  - heterogeneous SAN support, [210](#)
  - homogeneous SAN support, [209](#)
  - maximum configurations, [211](#)
  - product configurations, [209](#)
  - Small Business SAN, [209](#)
  - tape storage, [212](#)
  - zoning, [211](#)
- MSA1500
  - configuration rules, [208](#)
  - heterogeneous SAN support, [207](#)
  - homogeneous SAN support, [207](#)
  - implementation support, [207](#)
  - maximum configurations, [208](#)
  - tape storage, [209](#)
  - zoning, [208](#)
- MSA1510i
  - configuration rules, [318](#)
  - overview, [317](#)
  - support, [317](#)
- MSA2000i and MSA2000i G2, [318](#)
- multi-protocol long-distance technology, [259](#)
- multipath software
  - MPX200, [330](#)
- multipathing coexistence support
  - HP OpenVMS, [186](#)
  - HP Tru64 UNIX, [188](#)
  - HP-UX, [184](#)
  - IBM AIX, [191](#)
  - Microsoft Windows, [195](#)
  - Sun Solaris, [198](#)

## N

- N\_Port Virtualization mode, [165](#)
- N\_Port\_ID\_Virtualization
  - HBA NPIV, [170](#)
- NAS
  - configuration and zoning rules, [354](#)
  - HP ProLiant Storage Servers, [354](#)
  - SAN fabric rules, [356](#)
  - SAN integration, [354](#)
- network considerations, [261](#)
- nonrepudiation, [362](#)
- NonStop S-series servers support, [171](#)
- normalization, [261](#)
- NPIV
  - B-series switches, [89](#)
- NPV mode, [165](#)
  - configuration example, [166](#)
  - connectivity guidelines, [167](#)
  - considerations, [167](#)
  - failover policy, [166](#)
  - with FlexAttach, [168](#)
- NSPOF data protection, [40](#)
- NSPOF, EVA, [222](#)
- NX-OS, [77](#)

## O

- operating system

- zoning methods, 377
- operating system rules
  - B-series fabric, 95
  - C-series Fibre Channel switch fabric, 126
  - H-series switches, 135
  - heterogeneous, 160
  - HP-UX, 183
  - IBM AIX, 190
  - iSCSI Connectivity Option, 339
  - Linux, 192
  - M-series fabric, 143
  - Mac OS X, 188
  - Microsoft Windows, 193
  - OpenVMS, 185
  - Sun Solaris, 196
  - Tru64 UNIX, 187
  - VMware ESX, 198
- operating system support
  - FCoE C-series switches, 78
- optical add/drop multiplexers (OADMs), 254
- oversubscription, design considerations, 22
- overview
  - B-series switches, 83
  - C-series Fibre Channel and FCoE switches, 115
  - Fibre Channel routing, 43
  - H-series switches, 131
  - HP SVSP, 234
  - M-series switches, 140
  - SAN design, 16
  - SAN extension, 249
- P**
  - p-Class BladeSystem
    - Brocade 4 Gb SAN Switch, 91, 113
    - supported blade servers, 161
  - P2000 G3, 318
  - P6000 storage
    - configuration rules, 214
    - heterogeneous SAN support, 214
    - implementation support, 213
    - rules, 213
    - storage management server integration, 221
  - P6000 storage systems
    - H-series switches, 135
  - P6000/EVA connectivity
    - with MPX200, 329
  - P9000 storage systems
    - Continuous Access support, 296
  - P9000/XP storage systems
    - configuration rules, 227, 228
    - heterogeneous SAN support, 227
    - implementation support, 227
    - LUN Configuration and Security Manager, 231
    - tape storage system support, 229
    - zoning, 229
  - passive systems, 254
  - performance
    - guidelines, 157
    - maintaining beyond 5/10 km, 251

- recommendations, 156
- PID parameter bit, B-series switches, 112
- ports
  - covering, 374
  - setting, 10 to 100 km, 258
  - zoning method, 377

## Q

- Quality of Service, 262
- QuickLoop, interconnect, 146

## R

- recommendations
  - SAN best practices, 372
  - SAN designs, 20
- record keeping, connections, 374
- redundancy, routing, 51
- redundant active components
  - B-series Fibre Channel switches, 91
- redundant control processor
  - B-series Fibre Channel switches, 91
- regulations
  - security, 359
- related documentation, 393
- ring fabric
  - benefits, 29
  - expansion, 390
  - migrating from, 42
  - uses, 27
- ring fabric SAN
  - M-series switches maximum, 143
  - overview, 27
- routed fabric topologies, 35
- routed SAN fabric, 24
- router
  - connecting core-edge fabrics, 52
  - connecting fabrics through IP network, 53
- routing
  - comparing B-series, C-series, and H-series, 48
  - configurations, 52
  - features, 43
  - Fibre Channel implementations, 46
  - Fibre Channel integration FCIP, 57
  - methods, 48
  - overview, 43
  - SAN island consolidation, 57
  - SAN scaling, 57
  - summary, 57
  - table, 44
  - tape backup consolidation, 59
  - techniques, 44, 46
  - technology, 43
- rules
  - B-series fabric, 94
  - B-series Fibre Channel fabric, 95
  - C-series Fibre Channel switch fabric, 126
  - H-series switches, 135
  - interconnect, 146
  - iSCSI Connectivity Option, 340

- M-series fabric, 142
- MP Router fabrics, 107
- performance, 157
- switch, 20

## S

### SAN

- B-series addressing mode, 112
- bridging, 250
- components, 18
- configuration practices, 374
- congestion, 156
- data availability, 38
- design approaches, 20
- design specification, 372
- dual interoperable, 155
- expanding an existing topology, 389
- extending, 249
- heterogeneous interoperable, 154
- high bandwidth, 157
- high-availability dual redundant routed, 51
- homogeneous MSA1000, 209
- implementation practices, 372
- implementations, 17
- infrastructure, 19
- interoperable, 155
- latency, 156
- manageability, 22
- merging fabrics, 390
- mixed storage types, 200
- overview, 18
- P9000/XP/VA fabric boot support, 230
- planning, 372
- recommended practices, 372
- scaling, 19
- scaling and routing, 44
- security, 22
- solutions, 16
- SSP, 22
- standard designs, 20
- storage management server rules, 222
- topology design approaches, 23
- XP/VA shared fabric, 229

### SAN design

- configuration rules, 146
- interconnect rules, 146
- migration, 41
- overview, 16
- scalability, 41

### SAN extension

- best practices for HP P6000 Continuous Access, 289
- M-series, 258
- overview, 249
- products, 265
- types, 250

### SAN fabric

- design benefits, 24
- routed, 24
- topologies, 23

- zoning, 22

### SAN fabric topologies

- cascaded fabric SAN, 25
- congestion, 22
- core-edge, 29
- data access, 32
- design considerations, 20
- disaster tolerance, 21
- geographic layout, 20
- interoperability, 21
- layout, 20
- maximums, 33
- meshed fabric, 26
- migration, 21
- oversubscription, 20
- performance workload, 22
- ring fabric, 27
- scalability, 21
- single-switch, 24

### SAN Virtualization Services Platform, 234

### satellite switch

- support, 28

### scalability

- migration, 41
- MP Router, 110

### scaling

- nondisruptive methods, 389
- routing, 46
- switch, 44

### Secure Fabric OS

- B-series switches, 89

### security

- AES, 361
- attack classes, 358
- best practices, 358
- CHAP, 359
- compliance, 359
- confidentiality, 362
- data availability, 362
- data integrity, 362
- data protection, 362
- DES, 360
- encryption technologies, 360
- FC-SP, 360
- Fibre Channel technologies, 360
- HP Secure Advantage, 361
- international state regulations, 359
- IP SAN technologies, 359
- IPsec, 360
- iSCSI, 315
- key management, 361
- nonrepudiation, 362
- organizational policies, 361
- resource protection, 361
- technologies, 359
- types of threats, 358
- U.S. federal and state regulations, 359
- validation, 363

### Selective Storage Presentation, 22

- servers
  - blade support, 161
  - common access, 201
  - common HBA, 201
  - heterogeneous rules, 159
  - naming, 378
  - NonStop S-series support, 171
  - setup practices, 375
  - zoning rules, 202
- Service Location Protocol, 314
- SFP, 251
- shared link configuration, 260, 264
- single-switch fabric
  - benefits, 25
  - overview, 24
- skinny tree, 30
- small form factor pluggable see SFP
- soft zoning, discovery authentication, 376
- soft-plus zoning, login authentication, 376
- SONET, 264
- SSP
  - usage, 22
- storage
  - P6000 configuration rules, 213
- storage capacity, calculating, 21
- storage management server integration, 221, 246
- storage pools, 234
- storage product interface rules, 150
- storage replication products, 288
- storage security, 358
  - types of threats, 358
- storage system
  - HP 3PAR, 241
- storage systems
  - alias naming practices, 378
  - coexistence, 200
  - configuration practices, 375
  - ESX support, 198
  - EVA configuration rules, 213
  - HP OpenVMS support, 185
  - HP Tru64 UNIX support, 187
  - HP-UX support, 183
  - IBM AIX support, 190
  - Mac OS X support, 188
  - Microsoft Windows support, 193
  - MSA configuration rules, 203
  - P9000/XP configuration rules, 227
  - RA4000 configuration rules, 203
  - Sun Solaris support, 196
  - SVSP support, 237
  - VA configuration rules, 227
  - zoning by operating system, 376
- Subscriber's choice, HP, 393
- Sun Solaris
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 127
  - configuration rules, 197
  - M-series fabric rules, 143
  - multipathing coexistence support, 198
  - SAN rules, 196
  - supported storage systems, 196
- SVSP
  - Business Copy, 235
  - Continuous Access, 235
  - description, 234
  - features, 234
  - hardware requirements, 235
  - software, 235
  - storage pools, 234
  - storage rules, 237
  - thin provisioning, 234, 235
  - Volume Manager, 235
- switch models
  - B-series, 84
  - C-series FCoE expansion modules, 77
  - FCoE CN, 76
  - for cascaded fabric, 26
  - for core-edge fabric, 32
  - for meshed fabric, 27
  - for Meta SAN, 36
  - for ring fabric, 29
  - for single-switch fabric, 24
  - for VSAN, 37
  - H-series switches, 132
  - H-series switches with TR, 37
  - M-series, 140
  - McDATA, 140
  - MPX200 support, 330
- switch name, configuration settings, 391
- switch settings
  - B-series default, 94
  - H-series switches, 134
  - M-series default, 142
- switches
  - Blade server support, 161
  - mismatched parameters, 391
  - port covers, 374
  - third-party, 156
- symbols in text, 394
- Synchronous Optical Network, 264

## T

- tape backup, routing consolidation, 59
- tape storage systems
  - EVA, 221
  - H-series Fibre Channel Switch fabric rules, 135
  - MSA1000, 212
  - MSA1500, 209
  - P9500/XP/VA, 229
- technical support
  - HP, 393
- text symbols, 394
- ties, cable, 375
- time division multiplexing, 254
- top talkers
  - B-series switches, 89
- topologies
  - fabric types, 23

- SAN design approaches, 23
- SAN fabric, 23
- TR\_Port, 48
  - configuring, 56
- traffic isolation zones
  - B-series switches, 89
- transceiver
  - fiber optic, 251
  - GBIC, 251
- transparent router, 37
- transport distance rules, 150
- trunking
  - B-series switches, 88

## V

- VA storage systems
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 126
  - configuration rules, 227, 228
  - heterogeneous SAN support, 227
  - implementation support, 227
  - M-series fabric rules, 142
  - tape storage system support, 229
  - zoning, 229
- Virtual Connect
  - advantages, 162
  - configuration example, 162
  - guidelines and rules, 162
  - overview, 161
  - VC-FC, 161
- virtual disk, 234
- virtual fabrics
  - architecture, 48
  - B-series switches, 89
  - enabling B-series, 35
  - independence, 44
- virtual private networks, 262
- Virtualization Services Manager *see* VSM
- VMware ESX
  - C-series Fibre Channel switch fabric rules, 127
  - SAN rules, 198
  - supported storage systems, 198
- VSAN, 48, 50
  - benefits, 37
  - high-availability C-series fabric, 129
  - independence, 44
  - overview, 37
- VSM, 234

## W

- wavelength division multiplexing *see* WDM
- WDM
  - coarse wave division multiplexing, 254
  - network implementation, 253
  - system architectures, 253
  - system characteristics, 254
- websites
  - customer self repair, 395
  - HP Subscriber's choice for business, 393

- SAN Virtualization Services Platform, 235
- WWN
  - exporting, 44
  - zone member identifier, 378

## X

- Xen
  - configuration rules, 200
- XP storage system
  - with mpx110 rules, 275
- XP storage systems
  - B-series Fibre Channel switch fabric rules, 95
  - C-series Fibre Channel switch fabric rules, 127
  - M-series fabric rules, 143
- XPath OS, 110

## Z

- zone members
  - identifier types, 378
- zoning
  - B-series Fibre Channel switch guidelines, 113
  - B-series Fibre Channel switch limits, 112
  - C-series limits, 129
  - configuration mismatches, 391
  - enforcement, 376
  - EVA storage systems, 221
  - fabric, 22
  - H-series switches limits, 138
  - host based, 376
  - HP 3PAR, 246
  - M-series limits, 144
  - MSA1000, 211
  - MSA1500, 208
  - P9000/XP/VA storage systems, 229
- zoning enforcement
  - B-series Fibre Channel switches, 112
  - C-series Fibre Channel switches, 129
  - H-series switches, 138
  - M-series switches, 145
- zoning methods
  - application, 377
  - HBA, 377
  - operating system, 377
  - port allocation, 377